

TRILL Working Group
Internet Draft
Intended status: Informational

Tissa Senevirathne
Samer Salam
CISCO

Donald Eastlake
Sam Aldrin
HuaWei

Naveen Nimmu
Broadcom

Tal Mizrahi
Marvell

Anoop Ghanwani
Dell

June 25, 2012

Expires: December 2012

**Use of 802.1ag for TRILL OAM messages
draft-tissa-trill-8021ag-00.txt**

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), 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/ietf/1id-abstracts.txt>

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

This Internet-Draft will expire on December 25, 2012.

Copyright Notice

Copyright (c) 2012 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.

Abstract

In this document we present definitions of TRILL OAM messages. Messages defined in this document follow a similar structure to IEEE 802.1ag messages. In this document, only the high level proposal on using IEEE 802.1ag messaging is presented. The goal is to facilitate discussion on feasibility of using IEEE 802.1ag messages for TRILL OAM. Based on feedback from TRILL WG and IEEE 802.1 group, this document may be further updated.

Table of Contents

1.	Introduction.....	3
1.1.	Objective.....	4
2.	Conventions used in this document.....	4
3.	TRILL OAM Model.....	4
3.1	OAM Layering.....	4
3.2	TRILL OAM in RBridge Port Model.....	5
3.3	Maintenance Domains.....	6
3.4	MEPs and MIPs.....	7
4.	TRILL OAM Message channel.....	9
4.1.	TRILL OAM Message header.....	10
4.1.1.	Use of Magic Number and Checksum.....	11
4.2.	TRILL OAM Opcodes.....	11
4.3.	Format of TRILL OAM Message TLV.....	12
4.4.	TRILL OAM TLV.....	13
4.4.1.	Common TLV between 802.1ag and TRILL.....	13
4.4.2.	TRILL OAM TLV.....	13
4.4.2.1.	TRILL OAM Version TLV.....	13
4.4.3.	Out Of Band IP Address TLV.....	15
4.4.3.1.	Diagnostics VLAN TLV.....	15
4.4.3.2.	Original Data Payload TLV.....	16

5.1.1.	Loopback OAM Message format.....	17
5.1.2.	Theory of Operation.....	17
5.1.2.1.	Originator RBridge.....	17
5.1.2.2.	Intermediate RBridge.....	18
5.1.2.3.	Destination RBridge.....	18
5.2.	Loopback Message Hop-count method.....	19
5.2.1.	Prevent leaking out from TRILL network.....	19
6.	Path Trace Message.....	19
7.	Notification Messages.....	19
8.	Status Codes.....	20
8.1.	Error Codes.....	20
8.2.	Warning Codes.....	20
8.3.	Informational Codes.....	21
9.	Security Considerations.....	21
10.	IEEE Considerations.....	Error! Bookmark not defined.
11.	References.....	21
11.1.	Normative References.....	21
11.2.	Informative References.....	22
12.	Acknowledgments.....	22

1. Introduction

The structure of TRILL OAM messages is presented in [TRLOAMREQ]. Accordingly, TRILL OAM messages are constitute of four main parts; outer header, TRILL header, Flow Entropy and OAM message channel.

OAM Message channel allows defining various control information and carrying additional OAM related data between TRILL switches, also known as RBridges or Routing Bridges.

Outer header, TRILL header and Flow Entropy are technology (standard organization) dependent. OAM message channel, if defined properly, can be shared between different technologies. A common OAM channel allow a uniform user experience to the customers, savings on operator training, re-use of software code base and faster time to market.

In this document we propose to base the message channel on the [802.1ag] messaging format as defined in IEEE Connectivity Fault Management (CFM) [802.1ag] [802.1Q].

The ITU-T Y.1731 standard utilizes the same messaging format as [802.1ag]. However, IEEE defines a separate op-code space for the applicable messages specific to Y.1731. We propose a similar approach for TRILL and request a separate code space to be assigned for TRILL OAM messages.

1.1. Objective

The Objective of this document is to solicit feedback and comments on the use of 802.1ag messaging format as the OAM message format for TRILL. This document does not go into the operational details. Operational details are very similar to [[TLICMPOAM](#)]. Readers are referred to [[TLICMPOAM](#)] for details of operational aspects.

Updated and revised version of this document may be published in the future based on the feedback and comments of TRILL WG and IEEE 802.1 group.

2. Conventions used in this document

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](#)].

3. TRILL OAM Model

3.1. OAM Layering

In the RBridge architecture, the TRILL layer is independent of the underlying Link Layer technology. Therefore, it is possible to run TRILL over any transport layer capable of carrying Layer 2 frames such as Ethernet, PPP, or MPLS (Pseudo Wire). Furthermore, TRILL provides a virtual Ethernet connectivity service that is transparent to higher layer entities (e.g. Layer 3 and above). This strict layering is observed by TRILL OAM.

Of particular interest is the layering of TRILL OAM with respect to Ethernet CFM [[802.1ag](#)], especially that TRILL switches are likely to be deployed alongside existing 802.1 bridges in a network. Consider the example network depicted in Figure 1 below:

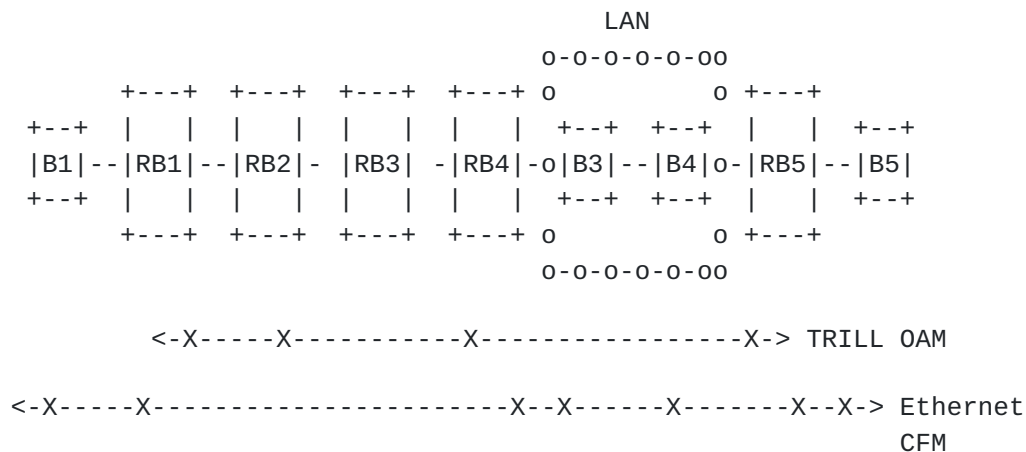


Figure 1: TRILL OAM & Ethernet CFM Layering

Where B_n and RB_n (n= 1,2,3,4,5) denote IEEE 802.1 bridges and TRILL Rbridges, respectively.

The "X" marks in the figure above indicate where each OAM protocol is applicable in the context of an example TRILL network. Note that this simplified diagram is not meant to capture the CFM Maintenance domain hierarchy nor the locality of MEPs/MIPs of various technologies. It is worth noting that an RBridge may actively participate in CFM only when connected to an 802.1 LAN. Transient Rbridges, in the core of a TRILL network, with no direct connectivity to 802.1 LAN are completely transparent to CFM.

3.2. TRILL OAM in RBridge Port Model

TRILL OAM processing can be modeled as a client of the Enhanced Internal Sublayer Service (EISS) in [802.1Q]. Furthermore, TRILL OAM requires services of the RBridge forwarding engine and utilizes information from the IS-IS control plane. Figure 2 below depicts TRILL OAM processing in the context of the RBridge port model defined in [RFC6325]. In this figure, double lines represent flow of both frames and information whereas single lines represent flow of information only.

While this figure shows a conceptual model, it is to be understood that implementations need not mirror this exact model as long as the intended OAM requirements and functionality are preserved.

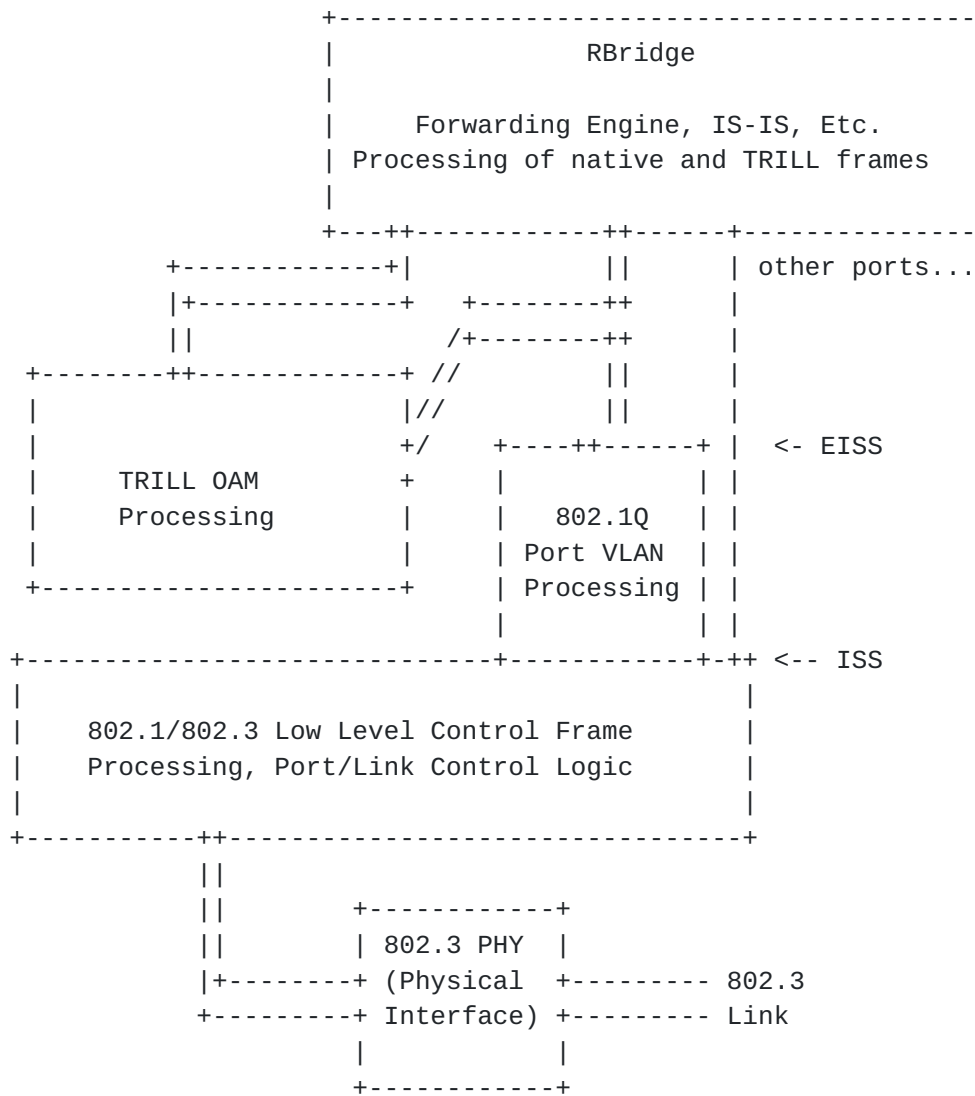


Figure 2: TRILL OAM in RBridge Port Model

3.3. Maintenance Domains

The concept of Maintenance Domains, or OAM Domains, is well known in the industry. IEEE 802.1ag, [RFC6136](#), [RFC5654](#), etc... all define the notion of a Maintenance Domain as a collection of devices (e.g. network elements) that are grouped for administrative and/or management purposes. Maintenance domains usually delineate trust

relationships, varying addressing schemes, network infrastructure capabilities, etc...

When mapped to TRILL, a Maintenance Domain is defined as a collection of RBridges in a network for which faults in connectivity or performance are to be managed by a single operator. All RBridges in a given Maintenance Domain are, by definition, owned and operated by a single entity (e.g. an enterprise or a data center operator, etc...). [RFC6325](#) defines the operation of TRILL in a single IS-IS area, with the assumption that the network is managed by a single operator. In this context, a single (default) Maintenance Domain is sufficient for TRILL OAM.

However, when considering scenarios where different TRILL networks need to be interconnected, for e.g. as discussed in [\[TRILLML\]](#), then the introduction of multiple Maintenance Domains and Maintenance Domain hierarchies becomes useful to map and contain administrative boundaries. When considering multi-domain scenarios, the following rules must be followed:

TRILL OAM domains MUST NOT overlap, but MUST either be disjoint or nest to form a hierarchy (i.e. a higher Maintenance Domain MAY completely include a lower Domain). A Maintenance Domain is typically identified by a Domain Name and a Maintenance Level (a numeric identifier). The larger the Domain, the higher the Level.

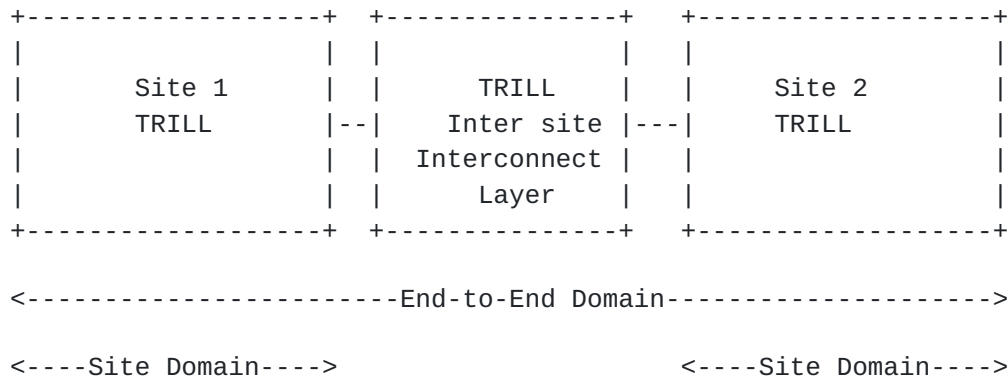


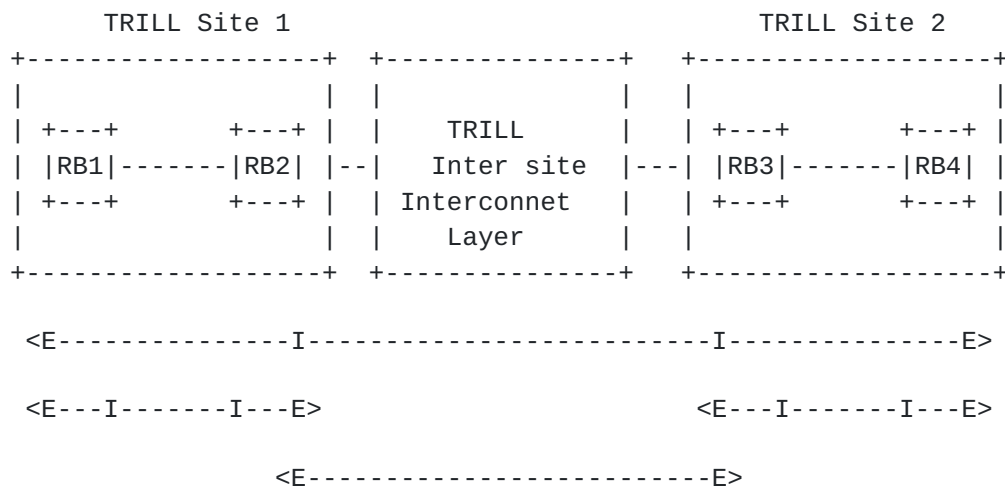
Figure 3: TRILL OAM Maintenance Domains

[3.4.](#) MEPS and MIPS

OAM capabilities on RBridges can be defined in terms of logical groupings of functions that can be categorized into two functional objects: TRILL Maintenance End Points (MEPs) and TRILL Maintenance Intermediate Points (MIPs).

MEPs are the active components of TRILL OAM: MEPs source TRILL OAM messages proactively or on-demand based on operator invocation. Furthermore, MEPs ensure that TRILL OAM messages do not leak outside the TRILL network and into end stations. MIPs, on the other hand, are internal to a Maintenance Domain. They are the passive components of TRILL OAM, primarily responsible for forwarding TRILL OAM messages and selectively responding to a subset of these messages.

The following figure shows the MEP and MIP placement for the Maintenance Domains depicted in Figure 3 above.



Legend E: MEP I: MIP

Figure 4: MEPs and MIPs

[802.1ag] specifies two distinct MP (Maintenance Points). Namely, Up MP and Down MP. [RFC6325] defines identification of TRILL frames received from the wire only. It does not define methods to identify frames egress in to the wire. Due to these reasons and to maintain simplicity, we propose only to define Down MP for TRILL OAM.

MEP/MIP of different technologies may exist on a given interface. Where applicable and Platforms MUST have capability to identify the applicable MIP/MEP.

3.5. TRILL OAM Maintenance Point (MP) Addressing

For unicast frames, TRILL MP is addressed by its TRILL nickname and either OAM Inner.MacSA or OAM Ethtype (Figure 5).

For multicast frames, TRILL MP is addressed by either Reserved Ethtype or Reserved source MAC (Figure 5).

For frames that include unmodified payloads, TRILL MP is addressed by its TRILL nickname, Hop-Count=0 (Figure 5). (NOTE: ingress RBridge set Hop-Count such that it count out at the intended RBridge). OAM signature allows OAM packets to be differentiated from data packets with Hop-Count=0

Following table summarizes the identification of different OAM frames from data frames.

Flow Entropy	OAM SRC MAC	OAM Eth Type	OAM Signature	RBridge nickname	Hop Count=0
unicast L2	N/A	Match	Optional	Match	N/A
Multicast L2	N/A	Match	Optional	N/A	N/A
Unicast IP	Match	N/A	Optional	Match	N/A
Multicast IP	Match	N/A	Optional	N/A	N/A
Notification	N/A	Match	Optional	Match	N/A
unmodified Payload	N/A	N/A	Match	Match	Match

Figure 5 Identification of OAM frames

4. TRILL OAM Message Channel

TRILL OAM Message Channel can be divided in to two parts. TRILL OAM Message header and TRILL OAM Message TLVs. Every OAM Message MUST contain a single TRILL OAM message header and set of one or more specified OAM Message TLVs.

4.1. TRILL OAM Message header

As discussed earlier, we propose to use the Message format defined in 802.1ag with some modifications. We believe these modifications facilitate a common messaging framework between [802.1ag], TRILL and other similar standards such as Y 1.731 and RFC6136

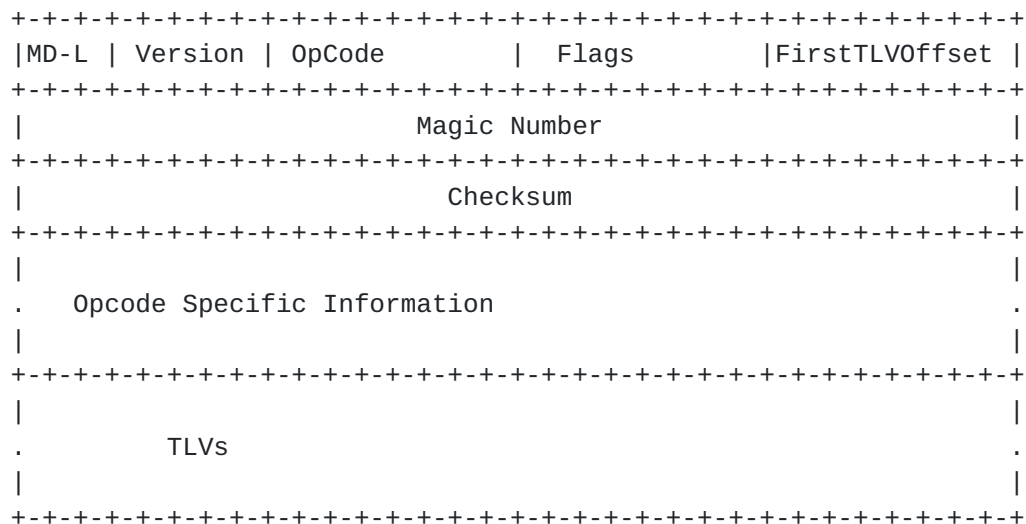


Figure 6 OAM Message Format

- o MD-L : Maintenance Domain Level (3 bits). Identifies the maintenance domain level. For TRILL this MAY be always set to zero. However, in MULTILEVEL TRILL, backbone MAY be of a different MD-LEVEL. (Please refer to [802.1ag] for the definition of MD-Level)
- o Version. Indicates the version (5 bits). [8021ag] set version to zero.
- o Flags: Include operational flags (1 byte). The definition of flags are opcode specific and is covered in the applicable sections.
- o FirstTLVOffset: Defines the location of the first TLV, in bytes, starting from the end of the FirstTLVOffset field. (1 byte) (Please refer to [802.1ag] for the definition of the FirstTLVOffset)
- o Magic Number: Increase the Entropy of the OAM checksum. Currently set to value (0x54729F74). (4 bytes).

Checksum : Calculated over MD-L, Version, OpCode, Flags, FirstTLVOffset. Checksum is calculated using FNV32 algorithm specified in [\[FNV\]](#). (4 bytes).

NOTE: Prior to calculating the checksum, implementations MAY pre validate the received frame by comparing the Version and Magic Number fields. Checksum is calculated only on frames that pass the pre-validation checks.

MD-L, Version, Opcode, Flags, FirstTLVOffset, Magicnumber and Checksum fields collectively are referred to as the OAM Message Header.

Opcode specific information section, based on the opcode, contains sequence number, time-stamp etc. Details about the Opcode specific information section and the associated TLVs are presented in other related documents.

4.1.1. Use of Magic Number and Checksum

The TRILL OAM framework requires the ability to differentiate between OAM packets and data packets. It is possible that applications may use real packets as the flow entropy. In such events, a reliable signature with a high degree of accuracy is required to differentiate between OAM and data packets. Version number is a 5 bit field and may not be adequate enough for this purpose. Implementations are required to first check for the matching Version number followed by the magic number. The checksum is calculated only if both the version and the checksum fields match the expected values. This procedure allows implementations (especially hardware) to execute checksum calculation process only on packets with higher probability of being an OAM packet. Packets with matching Version, Magic Number and Checksum fields are considered to be OAM packets.

As pointed out earlier and summarized in Figure 5, OAM signature validation is performed only on frames that were identified as OAM frames by the way of matching specific fields in the frame.

4.2. TRILL OAM Opcodes

We propose to define the following op-codes for TRILL. Each of the opcode defines a separate TRILL OAM message. Details of the messages are presented in the related sections. In this document we only define a selected few OAM messages. Based on the discussions and feedback, remaining OAM messages will be defined in future versions of this document.

TRILL OAM Message codes:

Senevirathne

Expires December 25, 2012

[Page 11]

```

64 : Loopback Message Reply
65 : Loopback Message
66 : Path Trace Reply
67 : Path Trace Message
68 : Notification Message
69 : Multicast Tree Verification Reply
70 : Multicast Tree Verification Message
71 : Loopback with Hop-count Reply
72 : Loopback with Hop-count Message.
73 : Performance Measurement one-way Reply
74 : Performance Measurement one-way Message
75 : Performance Measurement two-way Reply
76 : Performance Measurement two-way Message
77 - 95 : Reserved

```

4.3. Format of TRILL OAM Message TLV

We propose to use the same format as defined in section 21.5.1 of [802.1ag]. Following figure depicts the general format of TRILL OAM Message TLV.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |
|           Value(variable)         |
|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 7 TRILL OAM Message TLV

Type (1 octet) : Specifies the Type of the TLV (see sections 4.4.1. 4.4.2. for TLV types).

Length (2 octets) : Specifies the length of the values field in octets. Length of the value field can be either zero or more octets.

Value (variable): Length and the content of the value field depends on the type of the TLV. Please refer to applicable TLV definitions for the details.

Semantics and usage of Type values allocated for the TRILL OAM purpose are defined by this document and other future related documents.

4.4. TRILL OAM TLV

In this section we define the related TLVs. We propose to re-use [802.1ag] defined TLVs where applicable. Types 32-63 are reserved for ITUT Y.1731. We propose to reserve Types 64-95 for the purpose of TRILL OAM TLVs.

4.4.1. Common TLV between 802.1ag and TRILL

Following TLVs are defined in [802.1ag], we propose to re-use them where applicable. Format and semantics of the TLVs are as defined in [802.1g]. NOTE: Presented within brackets is the corresponding Type.

1. End TLV (0)
2. Sender ID TLV (1)
3. Port Status TLV (2)
4. Data TLV (3)
5. Interface Status TLV (4)
6. Reply Ingress TLV (5)
7. Reply Egress TLV (6)
8. LTM Egress Identifier TLV (7)
9. LTR Egress Identifier TLV (8)
10. Reserved (9-30)
11. Organization specific TLV (31)

4.4.2. TRILL OAM TLV

As indicated above, Types 64-95 will be requested to be reserved for TRILL OAM purpose. Listed below, a summary of TRILL OAM TLV and the corresponding codes. Format and semantics of TRILL OAM TLVs are defined in subsequent sections.

1. TRILL OAM Version (64)
2. Out of Band IP Address (65)
3. Diagnostic VLAN (66)
4. RBridge scope (67)
5. Original Payload TLV (68)
6. Reserved (69-95)

NOTE: Above is only for illustration purposes. In future revisions of this document, additional TLVs defined in [TLICMPOAM] will be defined within the above TLV space.

4.4.2.1. TRILL OAM Version TLV

TRILL OAM version can be different than the [802.1ag] version specified. TRILL OAM TLV specifies the TRILL OAM version. TRILL OAM TLV MUST be the first TLV in TRILL OAM messages. Messages that do

not include TRILL OAM TLV as the first TLV MUST be discarded.

Senevirathne

Expires December 25, 2012

[Page 13]


```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Type   | Length                               | Version   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Class|      Code      |      Reserved      |F|C|O|I|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 8 TRILL OAM Message TLV

Type (1 octet) = 64 indicate that this is the TRILL OAM Version

Length (2 octets) = 6

Version (1 Octet), currently set to zero. Indicates the TRILL OAM version. TRILL OAM version can be different that the [[802.1ag](#)] version.

Class (3 bits): Only applicable to Response and Notification messages.

0 : Success

1 : Error

2 : Warning

3 : Info

4 : 7 Reserved

Code (13 bits): Defined within the context of a class. Please see section Codes below for details.

Reserved: set to zero on transmission and ignored on reception.

F (1 bit) : Final flag, when set, indicates this is the last response.

C (1 bit): Cross connect error (VLAN mapping error), if set indicates VLAN cross connect error detected. This field is ignored in request messages and MUST only be interpreted in response messages.

O (1 bit) : If set, indicates, OAM out-of-band response requested.

I (1 bit) : If set, indicates, OAM in-band response requested.

NOTE: When both O and I bits are set to zero, indicates that no response is required. (silent mode). User MAY specify both O and I or one of them or none.

4.4.3. Out Of Band IP Address TLV

Out of Band IP Address TLV specifies the IP address to which out of band response MUST be sent. When 0 bit in the Version TLV is not set, Out of Band IP Address TLV is ignored. Length of the TLV implies the IP Address version.




Figure 9 Out of Band IP Address TLV

Type (1 octet) = 64 indicate that this is the TRILL OAM Version

Length (2 octets) = 5 or 17. Length Value 5 indicates it is IPv4 address and Length value of 17 indicates that it is IPv6 address.

IP Address (4 or 16 octets), valid IP address.

4.4.3.1. Diagnostics VLAN TLV

Diagnostic VLAN specifies the VLAN in which the OAM messages are generated on. Receiving RBridge MUST compare the inner.VLAN of the Flow entropy to the VLAN specified in the Diagnostic VLAN TLV. Cross connect Flag in the response MUST be set when the two VLANs do not match.




Figure 10 Diagnostic VLAN TLV

Type (1 octet) = 65 indicates that this is the TRILL Diagnostic VLAN TLV

Length (2 octets) = 5

V-Type (1 octet) 0- indicate 802.1Q 12 bit VLAN.

1 - indicate TRILL 24 bit fine grain label

Label (24 bits): Either 12 bit VLAN or 24 bit fine grain label.

4.4.3.2. Original Data Payload TLV




Figure 11 Out of Band IP Address TLV

Length (2 octets) = variable

5. Loopback Message

Loopback message is utilized for fault verification. It verifies connectivity between two RBridges, for a specified flow. Monitoring subsystem may use Loopback Message for connectivity monitoring and proactive fault detection. Users may specify exact flow, part of it or not at all. Additionally, users may also specify, ECMP choice at the ingress. ECMP choice can be a specific outgoing interface index, set of interface indexes, all of the interface indexes or none. If no ECMP index specified, payload is used to determine the ECMP choice. Method of deriving the ECMP choice using payload is implementation dependent and is outside the scope of this document. However, an implementation generating the Loopback message SHOULD use the same ECMP selection algorithm as the data plane. Additionally some implementation may allow users to specify the ingress interface that actual flow may ingress to the RBridge. Although ability to inject the data plane diagnostic frames from the ingress interface is an optional feature, it is highly desirable, as it allows verifying end-to-end connectivity from an ingress port to an egress RBridge.

Egress RBridge can send its response either in-band or out-of-band. In-band-response, additionally, allows measurement of the round trip delay. In-band responses are tagged with the same VLAN as the request frame. TRILL OAM TLVs in the request message allow the user to specify whether out-of-band response required. If out-of-band response is required, IP address to which the response is to be

directed MUST be specified.

Senevirathne Expires December 25, 2012

[Page 16]

Additionally, diagnostic VLAN, may be specified. Receiver RBridge may compare inner VLAN in the payload and the specified diagnostic VLAN. If the two specified VLAN values do not match, C flag in Version TLV SHOULD be set to indicate cross connect error.

5.1.1. Loopback OAM Message format

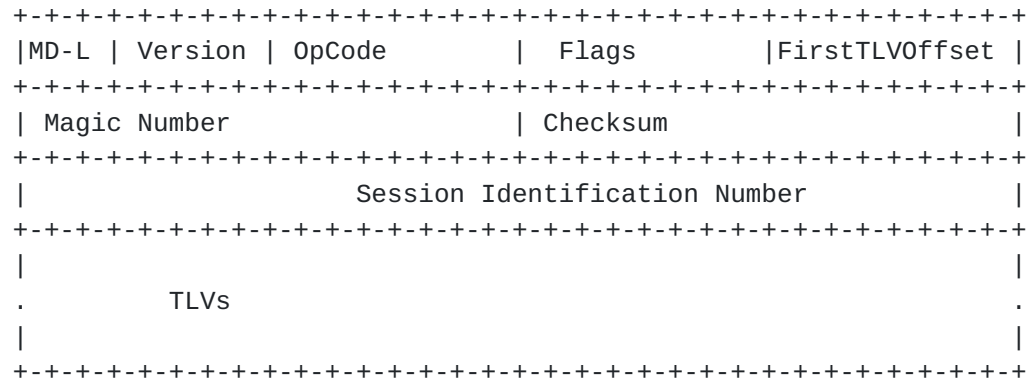


Figure 12 Loopback OAM Message Format

The above figure depicts the format of the Loopback Request and response messages. Opcode for Loopback Request Messages is set to 64 and Opcode of Response Message is set to 65. Session Identification Number is a 32 bit integer that allows the requesting RBridge to uniquely identify the corresponding session. Responding RBridges MUST echo the received "Session Identification" number without modifications

5.1.2. Theory of Operation

5.1.2.1. Originator RBridge

Identify the destination RBridge based on user specification or based on location of the specified address (see below sections for MAC discovery and address locator).

Construct the diagnostic payload based on user specified parameters. Default parameters MUST be utilized for unspecified payload parameters. See [[TLICMPOAM](#)] for default parameters.

Construct the TRILL OAM header. Set the opcode to (65), Loopback message. Assign applicable identification number and sequence number for the request.

TRILL OAM Version TLV MUST be included and set appropriate flags.

Senevirathne

Expires December 25, 2012

[Page 17]

Construct following ICMP multipart extensions, where applicable

- o Out-of-band response request
- o Out-of-band IP address
- o Diagnostic VLAN

Specify the Hop count of the TRILL data frame per user specification. Or utilize the applicable Hop count value, if TRILL TTL is not being specified.

Dispatch the diagnostic frame to the TRILL data plane for transmission.

RBridge may continue to retransmit the request at periodic intervals, until a response is received or the re-transmission count expires.

5.1.2.2. Intermediate RBridge

Intermediate RBridges forward the frame as a normal data frame and no special handling is required.

5.1.2.3. Destination RBridge

If the Loopback message is addressed to the local RBridge, then the RBridge process the message. The implementation performs frame validation and identifies OAM frames using methods specified in section 4.1.1. . The response is constructed as stated below.

Construction of the TRILL OAM response:

If out-of-band response is requested the destination IP address is set to the IP address specified in the request message. Source IP address is derived based on the outgoing IP interface address. UDP destination port is the TRILL OAM UDP port number.

Implementations encode the received TRILL header and flow entropy in the Original payload TLV.

If in-band response was requested, dispatch the frame to the TRILL data plane with request-originator RBridge nickname as the egress RBridge nickname.

If out-of-band response was requested, dispatch the frame to the standard IP forwarding process.

5.2. Loopback Message Hop-count method

The Loopback message procedures presented in [[TLICMPOAM](#)] Utilize customers specified payload to derive the diagnostic payload embedded in the OAM message.

From time to time, operators may desire the inclusion of actual user packets as the flow entropy of the OAM frame. The Hop-count method presented in this section facilitates inclusion of un-modified payload. When unmodified payloads are included as the diagnostic payload, there MUST be methods to identify such OAM frames from regular data frames and there MUST be methods to prevent such OAM frames from leaking out of TRILL network. Methods specified in [section 4.1.1](#). 4.1.1. allow R Bridges to identify OAM frames.

5.2.1. Prevent leaking out from TRILL network

First, the ingress RBridge that is generating the loopback message MUST discover the TRILL hop count to the egress RBridge. The discovered Hop count MUST be used as the hop count included in the TRILL header. Further, if the specified payload is IP, the IP header checksum SHOULD BE invalidated. The invalidation of IP checksum, prevents end stations further processing the OAM frames, in the unlikely event that it reached an end station due to a change in topology because of which the hop count is unable to suppress the message.

All other operations are similar to Loopback Message processing presented in [section 5.1.2](#).

6. Path Trace Message

Path Trace Message has the same format as Loopback Message. Opcode for Path Trace Request Message is 66 and Response 67.

Please refer to [[TLICMPOAM](#)] for Theory of Operation and details of Path Trace message.

7. Notification Messages

8.

TRILL OAM Notification message format is depicted in following figure.

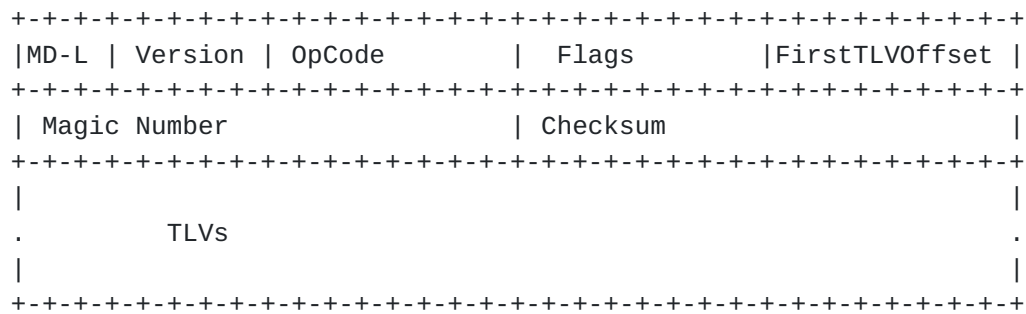


Figure 13 Notification OAM Message Format

The opcode of the Notification message is 68. Notification messages may be generated for variety of errors, warning and informational purposes. Notification messages are almost always asynchronous. Hence there is no Session Identification.

TRILL OAM Version TLV, which is mandatory, MUST be the first TLV. TRILL OAM Version TLV, has class field and code fields.

Class field specifies whether the message is Error, Warning or Informational Notification message. Code Field within the class specifies the actual notification.

9. Status Codes

Status codes are defined within a Class. Following Classes are currently defined.

Success, Error, Warning, Informational.

There are no status codes defined within the success class.

9.1. Error Codes

Following Error codes are currently defined

- 1: Egress RBridge Nickname unknown
- 2: Hop Count Expired
- 3: VLAN Unknown
- 4: Parameter Problem

9.2. Warning Codes

TBD

9.3. Informational Codes

TBD

10. Security Considerations

TBD

11. Allocation Considerations

10.1 IEEE Allocation Considerations

The IEEE 802.1 Working Group is requested to allocate a separate opcode and TLV space within 802.1g CFM messages for TRILL purpose.

10.2 IANA Considerations

- Set up sub-registry within the TRILL Parameters registry and initial entries for block of TRILL OAM OpCodes -
- Set up sub-registry within the TRILL Parameters registry and initial contents for TRILL OAM TLV Types -

12. References

12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [8021ag] IEEE, "Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management", 802.1ag, 2007.
- [8021Q] IEEE, "Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August 31st , 2011.
- [TLICMPOAM] Senevirathne, T., et.al., "ICMP based OAM Solution for TRILL", [draft-tissa-trill-oam-03](#), Work in Progress, January, 2012.
- [FNV] Fowler, G., et.al., "The FNV Non-Cryptographic Hash Algorithm", [draft-eastlake-fnv-03](#), Work in Progress, March, 2012.

12.2. Informative References

[RFC6325] Perlman, R., et.al., "Routing Bridges (RBridges): Base Protocol Specification", [RFC 6325](#), July 2011.

[TRILLML] Senevirathne, T., et.al., "Default Nickname Based Approach for Multi-level TRILL", [draft-tissa-trill-multilevel-00](#), Work in Progress, February 2012.

13. Acknowledgments

Work in this document was largely inspired by the directions provided by Stewart Bryant in finding common OAM solution between SD0.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Tissa Senevirathne
CISCO Systems
375 East Tasman Drive.
San Jose, CA 95134
USA.

Phone: +1 408-853-2291
Email: tsenevir@cisco.com

Samer Salam
CISCO Systems
595 Burrard St. Suite 2123
Vancouver, BC V7X 1J1, Canada

Email: ssalam@cisco.com

Donald Eastlake
Huawei Technologies
155 Beaver Street
Milford, MA 01757

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

Sam Aldrin
Huawei Technologies
2330 Central Express Way
Santa Clara, CA 95051
USA

Email: aldrin.ietf@gmail.com

Naveen Nimmu
Broadcom
9th Floor, Building no 9, Raheja Mind space
Hi-Tec City, Madhapur,
Hyderabad - 500 081, INDIA

Phone: +1-408-218-8893
Email: naveen@broadcom.com

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692 Israel

Email: talmi@marvell.com

Anoop Ghanwani
Dell
300 Holger Way,
San Jose, CA 95134, USA

Phone: +1-408-571-3500
Email: anoop@alumni.duke.edu.