

TRILL Working group
Internet Draft
Intended status: Standard Track
Updates: 6325

Tissa Senevirathne
Norman Finn
Samer Salam
Deepak Kumar
CISCO

Donald Eastlake
Sam Aldrin
Yizhou Li
Huawei

June 29, 2014

Expires: December 2014

TRILL Fault Management
draft-ietf-trill-oam-fm-06.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/lid-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 29, 2014.

Copyright Notice

Copyright (c) 2014 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

This document specifies TRILL OAM Fault Management. Methods in this document follow the IEEE 802.1 CFM (Continuity Fault Management) framework and reuse OAM tools where possible. Additional messages and TLVs are defined for TRILL specific applications or where a different set of information is required other than IEEE 802.1 CFM. This document updates RFC 6325.

Table of Contents

1. Introduction	4
2. Conventions used in this document	4
3. General Format of TRILL OAM Packets	5
3.1. Identification of TRILL OAM frames	7
3.2. Use of TRILL OAM Alert Flag	7
3.2.1. Handling of TRILL frames with the "A" Flag	8
3.3. OAM Capability Announcement	8
4. TRILL OAM Layering vs. IEEE Layering	9
4.1. Processing at ISS Layer	11
4.1.1. Receive Processing	11
4.1.2. Transmit Processing	11
4.2. End Station VLAN and Priority Processing	11
4.2.1. Receive Processing	11
4.2.2. Transmit Processing	11
4.3. TRILL Encapsulation and De-capsulation Layer	11
4.3.1. Receive Processing for Unicast packets	11
4.3.2. Transmit Processing for unicast packets	12
4.3.3. Receive Processing for Multicast packets	13
4.3.4. Transmit Processing of Multicast packets	14
4.4. TRILL OAM Layer Processing	15
5. Maintenance Associations (MA) in TRILL	16
6. MEP Addressing	17
6.1. Use of MIP in TRILL	20
7. Continuity Check Message (CCM)	22
8. TRILL OAM Message Channel	24

8.1.	TRILL OAM Message header	24
8.2.	TRILL Specific OAM Opcodes	25
8.3.	Format of TRILL OAM TLV	25
8.4.	TRILL OAM TLVs	26
8.4.1.	Common TLVs between CFM and TRILL	26
8.4.2.	TRILL OAM Specific TLVs	27
8.4.3.	TRILL OAM Application Identifier TLV	27
8.4.4.	Out Of Band Reply Address TLV	28
8.4.5.	Diagnostics Label TLV	29
8.4.6.	Original Data Payload TLV	30
8.4.7.	RBridge scope TLV	31
8.4.8.	Previous RBridge nickname TLV	32
8.4.9.	Next Hop RBridge List TLV	32
8.4.10.	Multicast Receiver Port count TLV	33
8.4.11.	Flow Identifier (flow-id) TLV	34
8.4.12.	Reflector Entropy TLV	34
8.4.13.	Authentication TLV	35
9.	Loopback Message	36
9.1.	Loopback OAM Message format	36
9.2.	Theory of Operation	37
9.2.1.	Actions by Originator RBridge	37
9.2.2.	Intermediate RBridge	37
9.2.3.	Destination RBridge	38
10.	Path Trace Message	38
10.1.	Theory of Operation	39
10.1.1.	Action by Originator RBridge	39
10.1.2.	Intermediate RBridge	40
10.1.3.	Destination RBridge	41
11.	Multi-Destination Tree Verification Message (MTVM)	41
11.1.	Multi-Destination Tree Verification Message (MTVM) Format	42
11.2.	Theory of Operation	42
11.2.1.	Actions by Originator RBridge	42
11.2.2.	Receiving RBridge	43
11.2.3.	In scope RBridges	43
12.	Application of Continuity Check Message (CCM) in TRILL ...	44
12.1.	CCM Error Notification	45
12.2.	Theory of Operation	46
12.2.1.	Actions by Originator RBridge	46
12.2.2.	Intermediate RBridge	47
12.2.3.	Destination RBridge	47
13.	Fragmented Reply	48
14.	Security Considerations	48
15.	IANA Considerations	49
15.1.	OAM Capabilitiy Flags	49
15.2.	CFM Code Points	50
15.3.	MAC Addresses	51

15.4. Return codes and sub codes	51
16. References	51
16.1. Normative References	51
16.2. Informative References	52
17. Acknowledgments	53
Appendix A. Backwards Compatibility	54
Appendix B. Base Mode for TRILL OAM	57
Appendix C. MAC Addresses Request	59

1. Introduction

The general structure of TRILL OAM messages is presented in [RFC7174]. TRILL OAM messages consist of five parts: link header, TRILL header, flow entropy, OAM message channel, and link trailer.

The OAM message channel carries various control information and OAM related data between TRILL switches, also known as RBridges or Routing Bridges.

A common OAM message channel representation can be shared between different technologies. This consistency between different OAM technologies promotes nested fault monitoring and isolation between technologies that share the same OAM framework.

The TRILL OAM message channel is formatted as specified in IEEE Connectivity Fault Management (CFM) [8021Q].

The ITU-T Y.1731 [Y1731] standard utilizes the same messaging format as [8021Q] OAM messages where applicable. This document takes a similar stance and reuses [8021Q] in TRILL OAM. It is assumed readers are familiar with [8021Q] and [Y1731]. Readers who are not familiar with these documents are encouraged to review them.

This document updates [RFC6325] as specified in Section 3.1.

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].

Acronyms used in the document include the following:

CCM - Continuity Check Message [8021Q]
ECMP - Equal Cost Multipath
ISS - Internal Sub Layer Service [8021Q]
LBM - Loop Back Message [8021Q]
LBR - Loop Back Reply Message [8021Q]
MP - Maintenance Point [RFC7174]
MEP - Maintenance End Point [RFC7174] [8021Q]
MIP - Maintenance Intermediate Point [RFC7174] [8021Q]
MA - Maintenance Association [8021Q] [RFC7174]
MD - Maintenance Domain [8021Q]
MTVM - Multi-destination Tree Verification Message
MTVR - Multi-destination Tree Verification Reply Message
OAM - Operations, Administration, and Maintenance [RFC6291]
PRI - Priority of Ethernet Frames [8021Q]
PTM - Path Trace Message
PTR - Path Trace Reply Message
TRILL - Transparent Interconnection of Lots of Links [RFC6325]
SAP - Service Access Point [8021Q]

3. General Format of TRILL OAM Packets

The TRILL forwarding paradigm allows an implementation to select a path from a set of equal cost paths to forward a unicast TRILL Data packet. For multi-destination TRILL Data packets, a distribution tree is chosen by the TRILL switch that ingresses or creates the packet. Selection of the path of choice is

implementation dependent at each hop for unicast and at the ingress for multi-destination. However, it is a common practice to utilize Layer 2 through Layer 4 information in the frame payload for path selection.

For accurate monitoring and/or diagnostics, OAM Messages are required to follow the same path as corresponding data packets. [RFC7174] presents the high-level format of the OAM messages. The details of the TRILL OAM frame format are defined in this document.

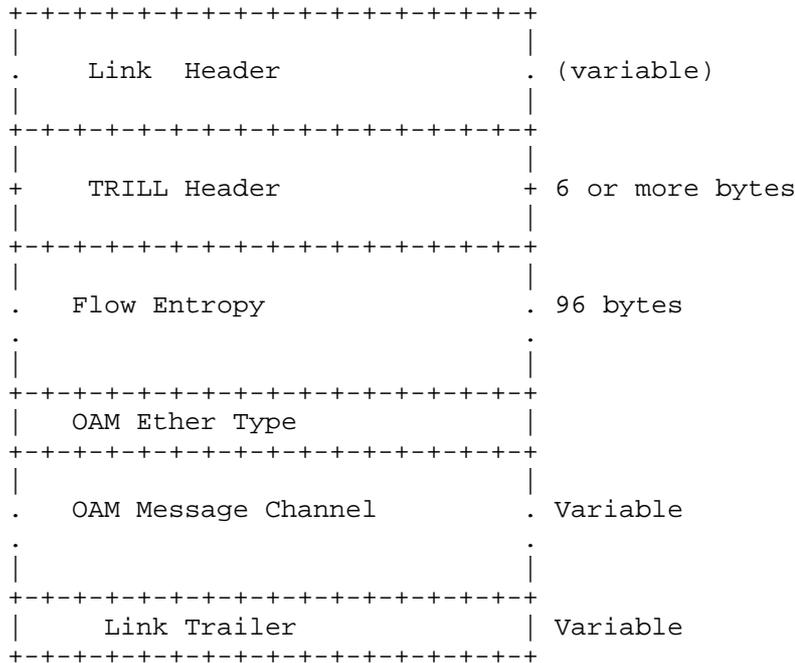


Figure 1 Format of TRILL OAM Messages

Link Header: Media-dependent header. For Ethernet, this includes Destination MAC, Source MAC, VLAN (optional) and EtherType fields.

TRILL Header: Fixed size of 6 bytes when the Extended Header is not included [RFC6325]

Flow Entropy: This is a 96-byte fixed size field. The rightmost bits of the field MUST be padded with zeros, up to 96 bytes, when the flow entropy is less than 96 bytes. Flow entropy enables emulation of the forwarding behavior of the desired data packets. The Flow Entropy field starts with the Inner.MacDA. The offset of the Inner.MacDA depends on whether extensions are included or not as specified in [RFC7179] and [RFC6325]. Such extensions are not commonly supported in current TRILL implementations.

OAM Ether Type: OAM Ether Type is 16-bit EtherType that identifies the OAM Message channel that follows. This document specifies using the EtherType 0x8902 allocated for CFM [8021Q]. OAM Message Channel: This is a variable size section that carries OAM related information. The message format is as specified in [8021Q].

Link Trailer: Media-dependent trailer. For Ethernet, this is the FCS (Frame Check Sequence).

3.1. Identification of TRILL OAM frames

TRILL, as originally specified in [RFC6325], did not have a specific flag or a method to identify OAM frames. This document updates [RFC6325] to include specific methods to identify TRILL OAM frames. Section 3.2. below explains the details of the method.

3.2. Use of TRILL OAM Alert Flag

The TRILL Header, as defined in [RFC6325], has two reserved bits. This document specifies use of the reserved bit next to Version field in the TRILL header as the Alert flag. Alert flag will be denoted by "A". Rbridges MUST NOT use the "A" flag for forwarding decisions such as the selection of which ECMP path or multi-destination tree to select.

Implementations that comply with this document MUST utilize "A" flag and CFM EtherType to identify TRILL OAM frames.

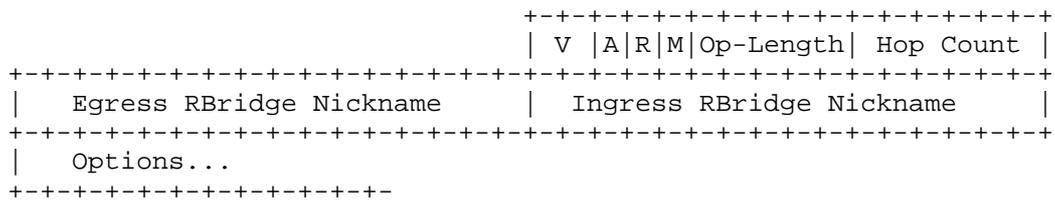


Figure 2 TRILL Header with the "A" Flag

A (1 bit) - Indicates this is a possible OAM frame and is subject to specific handling as specified in this document.

All other TRILL Header fields carry the same meaning as defined in RFC6325.

3.2.1. Handling of TRILL frames with the "A" Flag

Value "1" in the A flag indicates TRILL frames that may qualify as OAM frames. Implementations are further REQUIRED to validate such frames by comparing the value at the OAM Ether Type (Figure 1) location with the CFM EtherType "0x8902" [8021Q]. If the value matches, such frames are identified as TRILL OAM frames and SHOULD be processed as discussed in Section 4.

Frames with the "A" flag set that do not contain CFM EtherType are not considered as OAM frames. Such frames MUST be discarded.

3.3. OAM Capability Announcement

Any given RBridge can be (1) OAM incapable or (2) OAM capable with new extensions or (3) OAM capable with backwards-compatible method. The OAM request originator, prior to origination of the request is required to identify the OAM capability of the target and generate the appropriate OAM message.

Capability flags defined in TRILL version sub-TLV (TRILL-VER) [RFC7176] will be utilized for announcing OAM capabilities. The following OAM related capability flags are defined:

- O - OAM Capable
- B - Backwards Compatible OAM

A capability announcement, with "O" Flag set to 1 and "B" flag set to 1, indicates that the originating RBridge is OAM capable

but utilizes the backwards compatible method defined in Appendix A. A capability announcement with "O" Flag set to 1 and "B" flag set to 0, indicates that the originating RBridge is OAM capable and utilizes the method specified in section 3.2.

When "O" Flag is set to 0, the announcing implementation is considered not capable of OAM and the "B" flag is ignored.

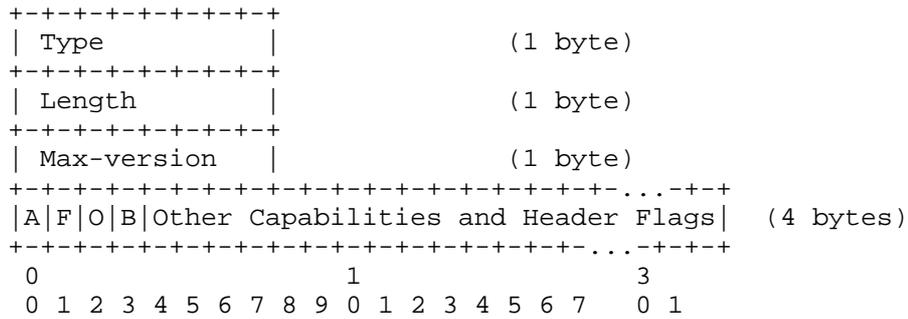


Figure 3 TRILL-VER sub-TLV [RFC7176] with O and B flags

Capability flags "A" and "F" are defined by [RFC7176] and [RFC7172]. "O" and "B" Flags are located after "F" flag in the Capability and Header Flags field of TRILL-VER sub-TLV, as depicted in Figure 3 above. Usage of "O" and "B" flags are as discussed above.

4. TRILL OAM Layering vs. IEEE Layering

This section presents the placement of the TRILL OAM shim within the IEEE 802.1 layers. The Transmit and Receive processing are explained.

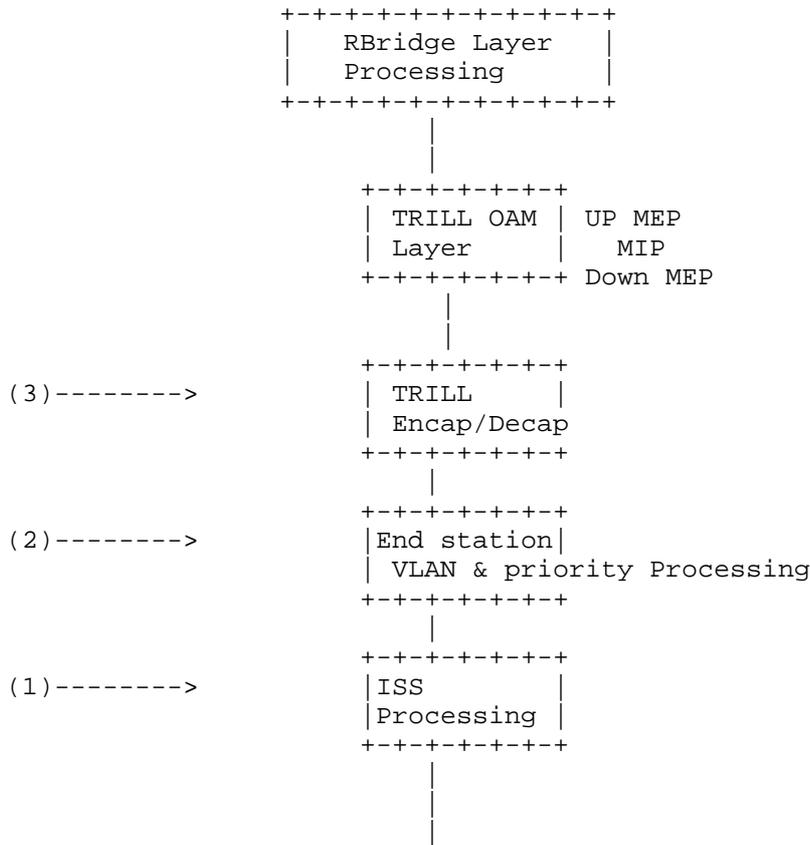


Figure 4 Placement of TRILL MP within IEEE 802.1

[RFC6325] Section 4.6 as updated by [RFC7180] provides a detailed explanation of frame processing. Please refer to those documents for additional details and for processing scenarios not covered herein.

Sections 4.1 and 4.2 below apply to links using a broadcast LAN technology such as Ethernet.

On links using an inherently point-to-point technology, such as PPP [RFC6361], there is no Outer.MacDA, Outer.MacSA, or Outer.VLAN because these are part of the link header for

Ethernet. Point-to-point links typically have link headers without these fields.

4.1. Processing at ISS Layer

4.1.1. Receive Processing

The ISS Layer receives an indication from the port. It extracts DA, SA and marks the remainder of the payload as M1. ISS Layer passes on (DA, SA, M1) as an indication to the higher layer.

For TRILL Ethernet frames, this is Outer.MacDA and Outer.MacSA. M1 is the remainder of the packet.

4.1.2. Transmit Processing

The ISS layer receives an indication from the higher layer that contains (DA, SA, M1). It constructs an Ethernet frame and passes down to the port.

4.2. End Station VLAN and Priority Processing

4.2.1. Receive Processing

Receives (DA, SA, M1) indication from ISS Layer. Extracts the VLAN ID and priority from the M1 part of the received indication (or derive them from the port defaults or other default parameters) and constructs (DA, SA, VLAN, PRI, M2). VLAN+PRI+M2 map to M1 in the received indication. Pass (DA, SA, VLAN, PRI, M2) to the TRILL encap/decap procession layer.

4.2.2. Transmit Procession

Receive (DA, SA, VLAN, PRI, M2) indication from TRILL encap/decap processing layer. Merge VLAN, PRI, M2 to form M1. Pass down (DA, SA, M1) to the ISS processing Layer.

4.3. TRILL Encapsulation and De-capsulation Layer

4.3.1. Receive Processing for Unicast packets

Receive indication (DA, SA, VLAN, PRI, M2) from End Station VLAN and Priority Processing Layer.

- o If DA matches port Local DA and Frame is of TRILL EtherType

- . Discard DA, SA, VLAN, PRI. From M2, derive (TRILL-HDR, iDA, iSA, i-VL, M3)
- . If TRILL nickname is Local and TRILL-OAM Flag is set
 - Pass on to OAM processing
- . Else pass on (TRILL-HDR, iDA, iSA, i-VL, M3) to RBridge Layer
 - o If DA matches port Local DA and EtherType is RBridge-Channel [RFC7178]
 - . Process as a possible unicast native RBridge Channel packet
 - o If DA matches port Local DA and EtherType is neither TRILL nor RBridge-Channel
 - . Discard packet
 - o If DA does not match and port is Appointed Forwarder for VLAN and EtherType is not TRILL or RBridge-Channel
 - . Insert TRILL-Hdr and send (TRILL-HDR, iDA, iSA, i-VL, M3) indication to RBridge Layer <- This is the TRILL Ingress Function.

4.3.2. Transmit Processing for unicast packets

- o Receive indication (TRILL-HDR, iDA, iSA, iVL, M3) from RBridge Layer
- o If egress TRILL nickname is local
 - o If port is Appointed Forwarder for iVL and the port is not configured as a trunk or p2p port and (TRILL Alert Flag set and OAM EtherType present) then
 - . Strip TRILL-HDR and construct (DA, SA, VLAN, M2) <- This is the TRILL Egress Function.
 - o Else
 - . Discard packet
- o If egress TRILL nickname is not local

- o Insert Outer.MacDA, Outer.MacSA, Outer.VLAN, TRILL EtherType and construct (DA, SA, VLAN, M2). Where M2 is (TRILL-HDR, iDA, iSA, iVL, M)
- o Forward (DA, SA, V, M2) to the VLAN End Station processing Layer.

4.3.3. Receive Processing for Multicast packets

- o Receive (DA, SA, V, M2) from VLAN aware end station processing layer
- o If the DA is All-RBridges and the EtherType is TRILL
 - o Strip DA, SA and V. From M2, extract (TRILL-HDR, iDA, iSA, iVL and M3).
 - o If TRILL Alert Flag is set and OAM EtherType is present at the end of Flow entropy
 - . Perform OAM Processing
 - o Else extract the TRILL header, inner MAC addresses and inner VLAN and pass indication (TRILL-HDR, iDA, iSA, iVL and M3) to TRILL RBridge Layer
- o If the DA is All-IS-IS-RBridges and the EtherType is L2-IS-IS then pass frame up to TRILL IS-IS processing
- o If the DA is All-RBridges or All-IS-IS-RBridges but EtherType is not TRILL or L2-IS-IS respectively
 - o Discard the packet
- o If the EtherType is TRILL but the multicast DA is not All-RBridges; or if the EtherType is L2-IS-IS but the multicast DA is not All-IS-IS-RBridges
 - o Discard the packet
- o If DA is All-Edge-RBridges and EtherType is RBridge-Channel [RFC7178]
 - o Process as a possible multicast native RBridge Channel packet

- o If the DA is in the initial bridging/link protocols block (01-80-C2-00-00-00 to 01-80-C2-00-00-0F) or is in the TRILL block and not assigned for Outer.MacDA use (01-80-C2-00-00-42 to 01-80-C2-00-00-4F) then
 - o The frame is not propagated through an RBridge although some special processing may be done at the port as specified in [RFC6325] and the frame may be dispatched to Layer 2 processing at the port if certain protocols are supported by that port (examples: Link Aggregation Protocol, Link Layer Discovery Protocol).
- o If the DA is some other multicast value
 - o Insert TRILL-HDR and construct (TRILL-HDR, iDA, iSA, iVL, M3)
 - o Pass the (TRILL-HDR, iDA, iSA, iVL, M3) to RBridge Layer

4.3.4. Transmit Processing of Multicast packets

The following ignores the case of transmitting TRILL IS-IS packets.

- o Receive indication (TRILL-HDR, iDA, iSA, iVL, M3) from RBridge layer.
- o If TRILL-HDR multicast flag set and TRILL-HDR Alert flag set and OAM EtherType present then:
 - o (DA, SA, V, M2) by inserting TRILL Outer.MacDA of All-RBridges, Outer.MacSA, Outer.VLAN and TRILL EtherType. M2 here is (EtherType TRILL, TRILL-HDR, iDA, iSA, iVL, M)
 - NOTE: Second copy of native format is not made.
- o Else If TRILL-HDR multicast flag set and Alert flag not set
 - o If the port is appointed Forwarder for iVL and the port is not configured as a trunk port or a p2p port, Strip TRILL-HDR, iSA, iDA, iVL and construct (DA, SA, V, M2) for native format.
 - o Make a second copy (DA, SA, V, M2) by inserting TRILL Outer.MacDA, Outer.MacSA, Outer.VLAN and TRILL

EtherType. M2 here is (EtherType TRILL, TRILL-HDR, iDA, iSA, iVL, M)

- o Pass the indication (DA, SA, V, M2) to End Station VLAN processing layer.

4.4. TRILL OAM Layer Processing

TRILL OAM Processing Layer is located between the TRILL Encapsulation / De-capsulation layer and RBridge Layer. It performs the following: 1. Identification of OAM frames that need local processing and 2. performs OAM processing or redirect to the CPU for OAM processing.

- o Receive indication (TRILL-HDR, iDA, iSA, iVL, M3) from RBridge layer. M3 is the payload after inner VLAN iVL.
- o If the TRILL Multicast Flag is set and TRILL Alert Flag is set and TRILL OAM EtherType is present then
 - o If MEP or MIP is configured on the Inner.VLAN/FGL of the packet then
 - . discard packets that have MD-LEVEL Less than that of the MEP or packets that do not have MD-LEVEL present (e.g., due to packet truncation).
 - . If MD-LEVEL matches MD-LEVEL of the MEP then
 - . Re-redirect to OAM Processing (Do not forward further)
 - . If MD-LEVEL matches MD-LEVEL of MIP then
 - . Make a Copy for OAM processing and continue
 - . If MD-LEVEL matches MD-LEVEL of MEP then
 - . Redirect the OAM packet to OAM processing and do not forward along or forward as a native packet.
 - o Else if TRILL Alert Flag is set and TRILL OAM EtherType is present then
 - o If MEP or MIP is configured on the Inner.VLAN/FGL of the packet then
 - . discard packets that have MD-LEVEL not present or MD-LEVEL is Less than that of the MEP.
 - . If MD-LEVEL matches MD-LEVEL of the MEP then
 - . Re-redirect to OAM Processing (Do not forward further)
 - . If MD-LEVEL matches MD-LEVEL of MIP then
 - . Make a Copy for OAM processing and continue

- o Else // Non-OAM Packet
 - o Continue
- o Pass the indication (DA, SA, V, M2) to End Station VLAN processing layer.

NOTE: In the Receive path, processing above compares against Down MEP and MIP Half functions. In the transmit processing it compares against Up MEP and MIP Half functions.

Appointed Forwarder is a function the TRILL Encap/De-Cap layer performs. The TRILL Encap/De-cap Layer is responsible for prevention of leaking of OAM packets as native frames.

5. Maintenance Associations (MA) in TRILL

[8021Q] defines a maintenance association as a logical relationship between a group of nodes. Each Maintenance Association (MA) is identified with a unique MAID of 48 bytes [8021Q]. CCM and other related OAM functions operate within the scope of an MA. The definition of MA is technology independent. Similarly it is encoded within the OAM message, not in the technology dependent portion of the packet. Hence the MAID as defined in [8021Q] can be utilized for TRILL OAM, without modifications. This also allows us to utilize CCM and LBM messages defined in [8021Q], as is.

In TRILL, an MA may contain two or more RBridges (MEPs). For unicast, it is likely that the MA contains exactly two MEPs that are the two end-points of the flow. For multicast, the MA may contain two or more MEPs.

For TRILL, in addition to all of the standard [8021Q] CFM MIB definitions, each MEP's MIB contains one or more flow entropy definitions corresponding to the set of flows that the MEP monitors.

[8021Q] CFM MIB is augmented to add the TRILL specific information. Figure 5, below depicts the augmentation of the CFM MIB to add the TRILL specific Flow Entropy.

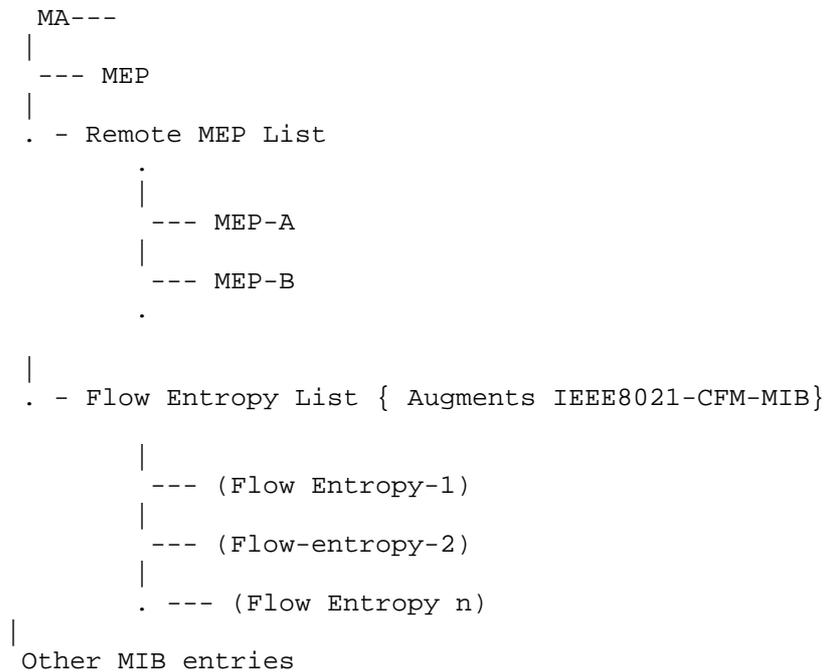


Figure 5 Correlation of TRILL augmented MIB

The detailed TRILL OAM MIB will be specified in a separate document [TRILLOAMMIB].

6. MEP Addressing

In IEEE CFM [8021Q], OAM messages address the target MEP by utilizing a unique MAC address. In TRILL a MEP is addressed by combination of the egress RBridge nickname and the Inner VLAN/FGL.

Additionally, MEPs are represented by 2 octet MEP-ID that is independent of the underlying technology. In CFM [8021Q] the value of MEP-ID is restricted to 1 to 8191. However, on CFM [8021Q] packet, MEP-ID are encoded as a 2 octet field. In TRILL Base Mode operation presented in Appendix B MEP-IDs are mapped 1

to 1 with the RBridge nicknames. Hence, In TRILL, MEP-ID MUST be a number in the range from 1 to 65535.

At the MEP, OAM packets go through a hierarchy of op-code de-multiplexers. The op-code de-multiplexers channel the incoming OAM packets to the appropriate message processor (e.g. LBM) The reader may refer to Figure 6 below for a visual depiction of these different de-multiplexers.

1. Identify the packets that need OAM processing at the Local RBridge as specified in Section 4.
 - a. Identify the MEP that is associated with the Inner.VLAN/FGL.
2. The MEP first validates the MD-LEVEL and then
 - a. Redirect to MD-LEVEL De-multiplexer
3. MD-LEVEL de-multiplexer compares the MD-Level of the packet against the MD level of the local MEPs of a given MD-Level on the port (Note: there can be more than one MEP at the same MD-Level but belonging to different MAs)
 - a. If the packet MD-LEVEL is equal to the configured MD-LEVEL of the MEP, then pass to the Opcode de-multiplexer
 - b. If the packet MD-LEVEL is less than the configured MD-LEVEL of the MEP, discard the packet
 - c. If the packer MD-LEVEL is greater than the configured MD-LEVEL of the MEP, then pass on to the next higher MD-LEVEL de-multiplexer, if available. Otherwise, if no such higher MD-LEVEL de-multiplexer exists, then forward the packet as normal data.
4. Opcode De-multiplexer compares the opcode in the packet with supported opcodes
 - a. If Op-code is CCM, LBM, LBR, PTM, PTR, MTVM, MTRV, then pass on to the correct Processor
 - b. If Op-code is Unknown, then discard.

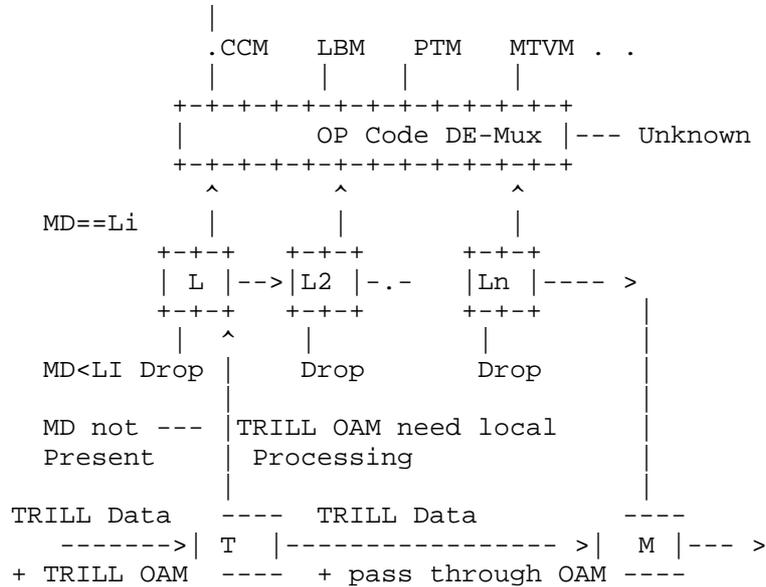


Figure 6 OAM De-Multiplexers at MEP for active SAP

T : Denotes Tap, that identifies OAM frames that need local processing. These are the packets with Alert flag set and OAM EtherType is present after the flow entropy of the packet

M : Is the post processing merge, merges data and OAM messages that are passed through. Additionally, the Merge component ensures, as explained earlier, that OAM packets are not forwarded out as native frames.

L : Denotes MD-Level processing. Packets with MD-Level less than the Level will be dropped. Packets with equal MD-Level are passed on to the opcode de-multiplexer. Others are passed on to the next level MD processors or eventually to the merge point (M).

NOTE: LBM, LBR, MTVM, MTRV, PTM and PTR are not subject to MA de-multiplexers. These packets do not have an MA encoded in the packet. Adequate response can be generated to these packets, without loss of functionality, by any of the MEPs present on that interface or an entity within the RBridge.

6.1. Use of MIP in TRILL

Maintenance Intermediate Points (MIP) are mainly used for fault isolation. Link Trace Messages in [8021Q] utilize a well-known multicast MAC address and MIPs generate responses to Link Trace messages. Response to Link Trace messages or lack thereof can be used for fault isolation in TRILL.

As explained in section 10. , a hop-count expiry approach will be utilized for fault isolation and path tracing. The approach is very similar to the well-known IP trace-route approach. Hence, explicit addressing of MIPs is not required for the purpose of fault isolation.

Any given RBridge can have multiple MIPs located within an interface. As such, a mechanism is required to identify which MIP should respond to an incoming OAM message.

A similar approach to that presented above for MEPs can be used for MIP processing. It is important to note that "M", the merge block of a MIP, does not prevent OAM packets leaking out as native frames. On edge interfaces, MEPs MUST be configured to prevent the leaking of TRILL OAM packets out of the TRILL Campus.

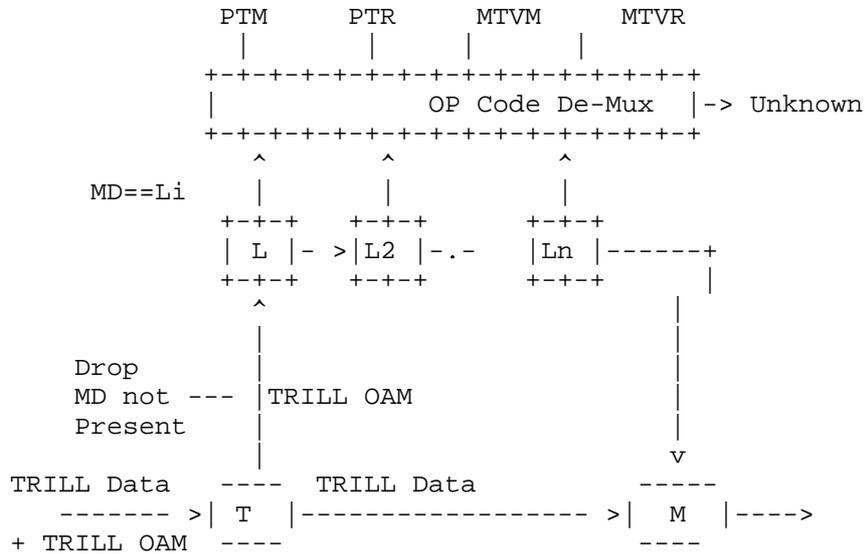


Figure 7 OAM De-Multiplexers at MIP for active SAP

T: TAP processing for MIP. All packets with OAM flag set are captured.

L : MD Level Processing, Packet with matching MD Level are "copied" to the Opcode de-multiplexer and original packet is passed on to the next MD level processor. Other packets are simply passed on to the next MD level processor, without copying to the OP code de-multiplexer.

M : Merge processor, merge OAM packets to be forwarded along with the data flow.

Packets that carry Path Trace Message (PT) or Multi-destination Tree Verification (MTVM) OpCodes are passed on to the respective processors.

Packets with unknown OpCodes are counted and discarded.

7. Continuity Check Message (CCM)

CCMs are used to monitor connectivity and configuration errors. [8021Q] monitors connectivity by listening to periodic CCM messages received from its remote MEP partners in the MA. An [8021Q] MEP identifies cross-connect errors by comparing the MAID in the received CCM message with the MEP's local MAID. The MAID [8021Q] is a 48-byte field that is technology independent. Similarly, the MEPID is a 2-byte field that is independent of the technology. Given this generic definition of CCM fields, CCM as defined in [8021Q] can be utilized in TRILL with no changes. TRILL specific information may be carried in CCMs when encoded using TRILL specific TLVs or sub-TLVs. This is possible since CCMs may carry optional TLVs.

Unlike classical Ethernet environments, TRILL contains multipath forwarding. The path taken by a packet depends on the payload of the packet. The Maintenance Association identifies the interested end-points (MEPs) of a given monitored path. For unicast there are only two MEPs per MA. For multicast there can be two or more MEPs in the MA. The entropy values of the monitored flows are defined within the MA. CCM transmit logic will utilize these flow entropy values when constructing the CCM packets. Please see section 12. below for the theory of operation of CCM.

The MIB of [8021Q] is augmented with the definition of flow-entropy. Please see [TRILLOAMMIB] for definition of these and other TRILL related OAM MIB definitions. The below Figure depicts the correlation between MA, CCM and the flow-entropy.

```

    MA---
    |
    --- MEP
    |
    . - Remote MEP List
        .
        |
        --- MEP-A
        |
        --- MEP-B
        .

    |
    . - Flow Entropy List {Augments IEEE8021-CFM-MIB}
        |
        --- (Flow Entropy-1)
        |
        --- (Flow-entropy-2)
        |
        . ---(Flow Entropy n)

    |
    . - CCM
        |
        --- (standard 802lag entries)
        |
        --- (hop-count) { Augments IEEE8021-CFM-MIB}
        |
        --- (Other TBD TRILL OAM specific entries)
                                {Augmented}

    |
    .
    |
    - Other MIB entries

```

Figure 8 Augmentation of CCM MIB in TRILL

In a multi-pathing environment, a Flow - by definition - is unidirectional. A question may arise as to what flow entropy should be used in the response. CCMs are unidirectional and have no explicit reply; as such, the issue of the response flow entropy does not arise. In the transmitted CCM, each MEP reports local status using the Remote Defect Indication (RDI) flag. Additionally, a MEP may raise SNMP TRAPS [TRILLOAMMIB] as Alarms when a connectivity failure occurs.

8. TRILL OAM Message Channel

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

8.1. TRILL OAM Message header

As discussed earlier, a common messaging framework between [8021Q], TRILL, and other similar standards such as Y.1731 is accomplished by re-using the OAM message header defined in [8021Q].

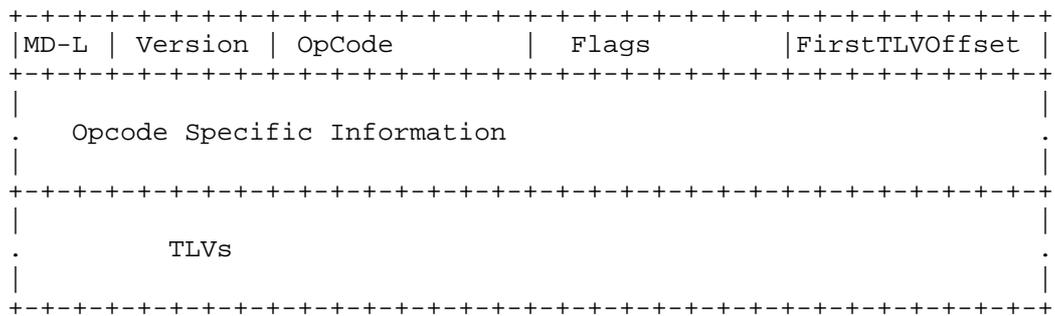


Figure 9 OAM Message Format

- o MD-L: Maintenance Domain Level (3 bits). Identifies the maintenance domain level. For TRILL, in general, this field is set to a single value across the TRILL campus. When using TRILL base mode as specified in Appendix B, MD-L is set to 3. However, extension of TRILL, for example to support multilevel, may create different MD-LEVELs and MD-L field must be appropriately set in those scenarios. (Please refer to [8021Q] for the definition of MD-Level)
- o Version: Indicates the version (5 bits) as specified in [8021Q]. This document does not require changing the Version defined in [8021Q].
- o Flags: Includes operational flags (1 byte). The definition of flags is 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). (Refer to [8021Q] for the definition of the FirstTLVOffset.)

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

The Opcode specific information section of the OAM Message may contain Session Identification number, time-stamp, etc.

8.2. TRILL Specific OAM Opcodes

The following TRILL specific CFM Opcodes are defined. Each of the Opcodes indicates a separate type of TRILL OAM message. Details of the messages are presented in the related sections.

TRILL OAM Message Opcodes:

- TBD1: Path Trace Reply
- TBD2: Path Trace Message
- TBD3: Multicast Tree Verification Reply
- TBD4: Multicast Tree Verification Message

Loopback and CCM Messages reuse the opcodes defined by [8021Q]

8.3. Format of TRILL OAM TLV

The same CFM TLV format as defined in [8021Q] is used for TRILL OAM. The following figure depicts the general format of a TRILL OAM TLV:

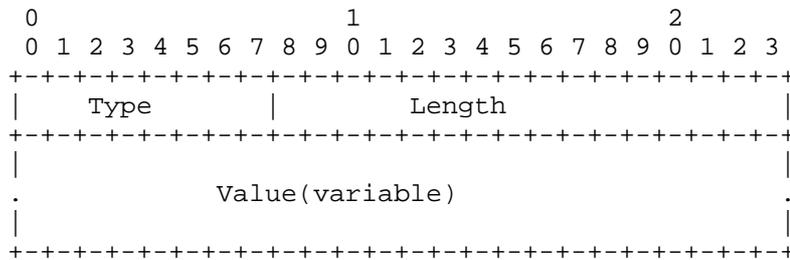


Figure 10 TRILL OAM TLV

Type (1 octet): Specifies the Type of the TLV (see sections 8.4. for TLV types).

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

Value (variable): The length and the content of this field depend on the type of the TLV. Please refer to applicable TLV definitions for the details.

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

8.4. TRILL OAM TLVs

TRILL related TLVs are defined in this section. [8021Q] defined TLVs are reused, where applicable.

8.4.1. Common TLVs between CFM and TRILL

The following TLVs are defined in [8021Q]. We re-use them where applicable. The format and semantics of the TLVs are as defined in [8021Q].

Type	Name of TLV in [8021Q]
----	-----
0	End TLV
1	Sender ID TLV
2	Port Status TLV
3	Data TLV
4	Interface Status TLV
5	Reply Ingress TLV
6	Reply Egress TLV
7	LTM Egress Identifier TLV
8	LTR Egress Identifier TLV
9-30	Reserved
31	Organization Specific TLV

8.4.2. TRILL OAM Specific TLVs

Listed below is a summary of TRILL OAM TLVs and their corresponding codes. Format and semantics of TRILL OAM TLVs are defined in subsequent sections.

Type	TLV Name
TBDa	TRILL OAM Application Identifier TLV
TBDb	Out of Band Reply Address TLV
TBDc	Diagnostic Label TLV
TBDd	Original Data Payload TLV
TBDe	RBridge scope TLV
TBDf	Previous RBridge nickname TLV
TBDg	Next Hop RBridge List (ECMP) TLV
TBDh	Multicast Receiver Port count TLV
TBDi	Flow Identifier TLV
TBDj	Reflector Entropy TLV
TBDk	Authentication TLV

8.4.3. TRILL OAM Application Identifier TLV

TRILL OAM Application Identifier TLV carries TRILL OAM application specific information. The TRILL OAM Application Identifier TLV MUST always be present and MUST be the first TLV in TRILL OAM messages. Messages that do not include the TRILL OAM Application Identifier TLV as the first TLV MUST be discarded by TRILL MP.

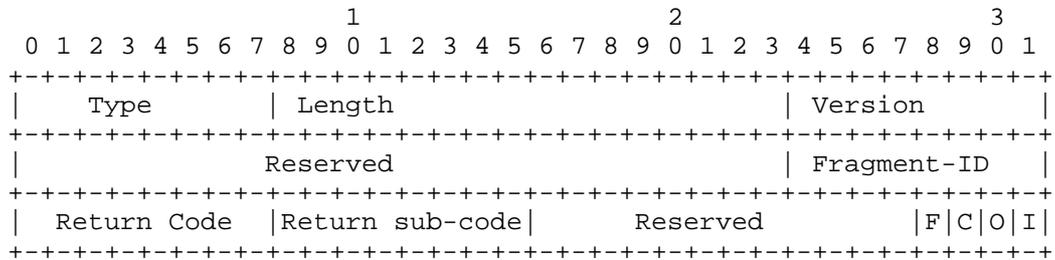


Figure 11 TRILL OAM Application Identifier TLV

Type (1 octet) = TBDa indicate that this is the TRILL OAM Application Identifier TLV

Length (2 octets) = 6

TRILL OAM Version (1 Octet), currently set to zero. Indicates the TRILL OAM version. TRILL OAM version can be different than the [8021Q] version.

Fragment-ID (1 octet): Indicates the fragment number of the current message. This applies only to the reply messages. F flag defined below MUST be set with the final message whether it is the last fragment of the fragmented message or only message of the reply. Section 13. below provide more details.

Return Code (1 Octet): Set to zero on requests. Set to an appropriate value in response messages.

Return sub-code (1 Octet): Return sub-code is set to zero on transmission of request message. Return sub-code identifies categories within a specific Return code. Return sub-code MUST be interpreted within a Return code.

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 flag(VLAN/Label mapping error), if set indicates that the label (VLAN/FGL) in the flow entropy is different than the label included in the diagnostic TLV. 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. When both O and I bits are set response is sent both in-band and out-of-band.

8.4.4. Out Of Band Reply Address TLV

Out of Band Reply Address TLV specifies the address to which an out of band OAM reply message MUST be sent. When O bit in the TRILL Version TLV is not set, Out of Band Reply Address TLV is ignored.

Type (1 Octet) = TBDd

Length (2 octets) = variable

Original Payload: The original TRILL Header and Entropy. Used in constructing replies to the Loopback Message (see Section 9) and the Path Trace Message (see Section 10).

8.4.7. RBridge scope TLV

RBridge scope TLV identifies nicknames of RBridges from which a response is required. The RBridge scope TLV is only applicable to Multicast Tree Verification messages. This TLV SHOULD NOT be included in other messages. Receiving RBridges MUST ignore this TLV on messages other than Multicast Verification Message.

Each TLV can contain up to 255 nicknames of in-scope RBridges. A Multicast Verification Message may contain multiple "RBridge scope TLVs", in the event that more than 255 in scope RBridges need to be specified.

Absence of the "RBridge scope TLV" indicates that a response is needed from all the RBridges. Please see section 11. for details.

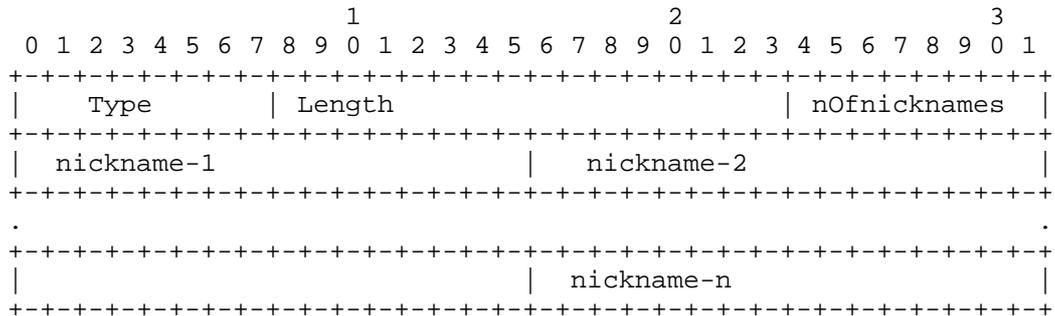


Figure 15 RBridge Scope TLV

Type (1 octet) = TBDe indicates that this is the "RBridge scope TLV"

Length (2 octets) = variable. Minimum value is 2.

Nickname (2 octets) = 16 bit RBridge nickname.

8.4.8. Previous RBridge nickname TLV

"Previous RBridge nickname TLV" identifies the nickname or nicknames of the Previous RBridge. [RFC6325] allows a given RBridge to hold multiple nicknames.

"Previous RBridge nickname TLV" is an optional TLV. Multiple instances of this TLV MAY be included when an upstream RBridge is represented by more than 255 nicknames (highly unlikely).

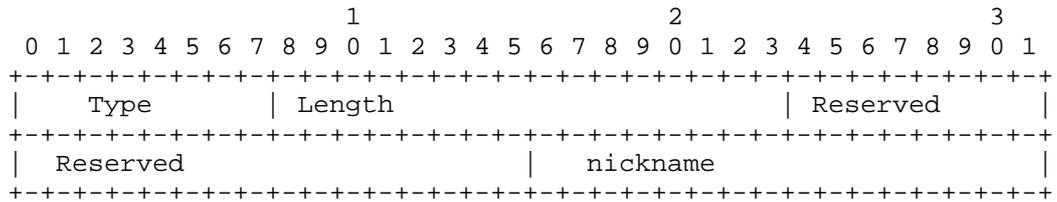


Figure 16 Previous RBridge nickname TLV

Type (1 octet) = TBDf indicates that this is the "Previous RBridge nickname"

Length (2 octets) = 4.

Nickname (2 octets) = 16 bit RBridge nickname.

8.4.9. Next Hop RBridge List TLV

"Next Hop RBridge List TLV" identifies the nickname or nicknames of the downstream next hop RBridges. [RFC6325] allows a given RBridge to have multiple Equal Cost Paths to a specified destination. Each next hop RBridge is represented by one of its nicknames.

"Next Hop RBridge List TLV" is an optional TLV. Multiple instances of this TLV MAY be included when there are more than 255 Equal Cost Paths to the destination.

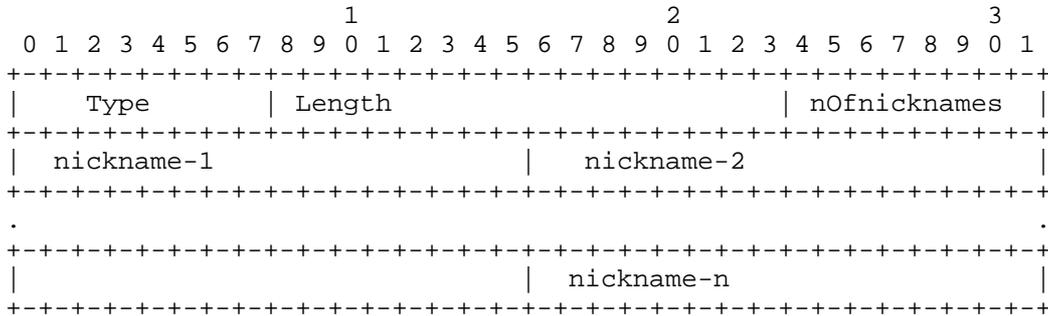


Figure 17 Next Hop RBridge List TLV

Type (1 octet) = TBDg indicates that this is the "Next nickname"

Length (2 octets) = variable. Minimum value is 2.

Nickname (2 octets) = 16 bit RBridge nickname.

8.4.10. Multicast Receiver Port count TLV

"Multicast Receiver Port Count TLV" identifies the number of ports interested in receiving the specified multicast stream within the responding RBridge on the label (VLAN or FGL) specified by the Diagnostic Label TLV.

Multicast Receiver Port count is an Optional TLV.

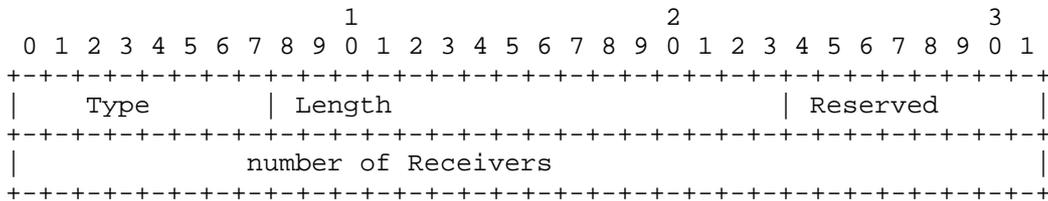


Figure 18 Multicast Receiver Availability TLV

Type (1 octet) = TBDh indicates that this is the "Multicast Availability TLV"

Length (2 octets) = 5.

Number of Receivers (4 octets) = Indicates the number of Multicast receivers available on the responding RBridge on the label specified by the diagnostic label.

8.4.11. Flow Identifier (flow-id) TLV

Flow Identifier (flow-id) uniquely identifies a specific flow. The flow-id value is unique per MEP and needs to be interpreted as such.

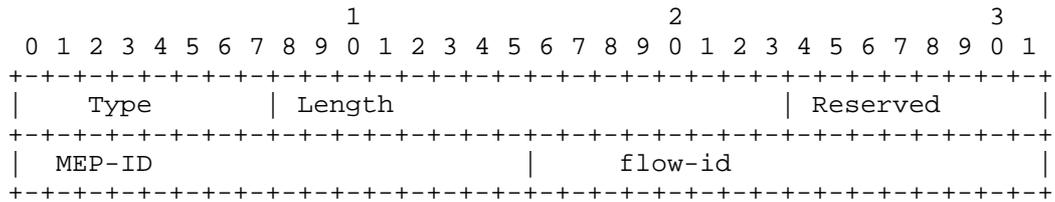


Figure 19 Flow Identifier TLV

Type (1 octet) = TBDi

Length (2 octets) = 5.

Reserved (1 octet) set to 0 on transmission and ignored on reception.

MEP-ID (2 octets) = MEP-ID of the originator [8021Q]. In TRILL MEP-ID can take a value from 1 to 65535.

Flow-id (2 octets) = uniquely identifies the flow per MEP. Different MEPs may allocate the same flow-id value. The {MEP-ID, flow-id} pair is globally unique.

Inclusion of the MEP-ID in the flow-id TLV allows the inclusion of a MEP-ID for messages that do not contain a MEP-ID in their OAM header. Applications may use MEP-ID information for different types of troubleshooting.

8.4.12. Reflector Entropy TLV

Reflector Entropy TLV is an optional TLV. This TLV, when present, tells the responder to utilize the Reflector Entropy specified within the TLV as the flow-entropy of the response message.

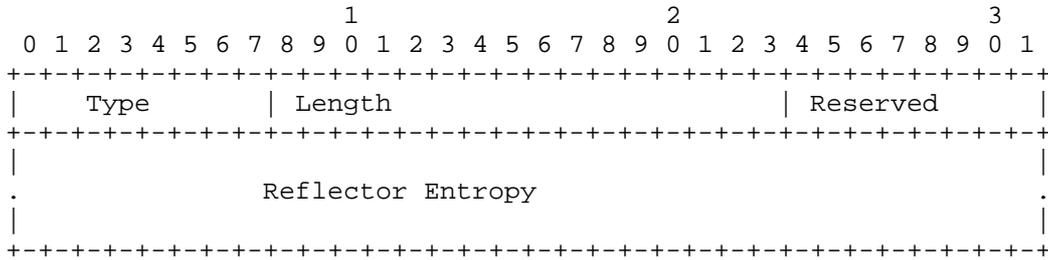


Figure 20 Reflector Entropy TLV

Type (1 octet) =TBDj Reflector Entropy TLV.

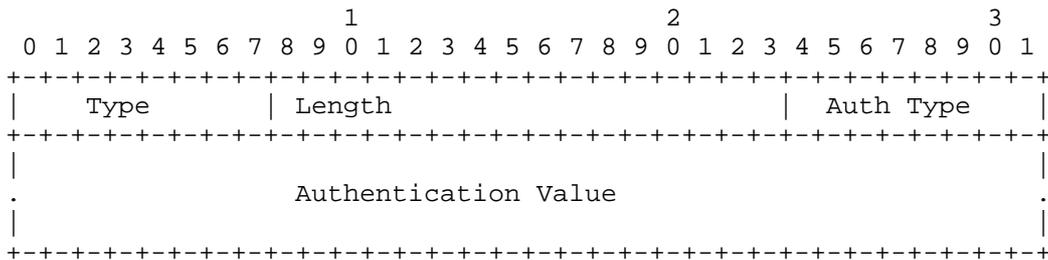
Length (1 octet) =97.

Reserved (1 octet) = set to zero on transmission and ignored by the recipient.

Reflector Entropy (96-octet) = Flow Entropy to be used by the responder. May be padded with zero if the desired flow entropy is less than 96 octets.

8.4.13. Authentication TLV

The Authentication TLV is an optional TLV that can appear in any OAM Message or Reply in TRILL.



Type (1 octet) =TBDk Authentication TLV.

Length (1 octet) = variable length

The Auth Type and following Authentication Value are the same as the Auth Type and following value for the [IS-IS] Authentication TLV. It is RECOMMENDED that Auth Type 3 be used, in which case the process is generally as specified in [RFC5310] using the same Key ID space as TRILL [IS-IS]. The area covered by the Authentication TLV is from the beginning of the TRILL Header to the end of the TRILL OAM Message Channel - the Link Header and Trailer are not included. The TRILL Header Alert and Reserved bit and Hop Count are treated as if zero for the purposes of computing and verifying the Authentication Data.

An RBridge supporting OAM authentication can be configured to either (1) ignore received OAM Authentication TLVs and not send them, (2) ignore received OAM Authentication TLVs but include them in all OAM packets sent, or (3) to include Authentication TLVs in all OAM messages sent and enforce authentication of OAM messages received. When an RBridge is enforcing authentication, it discards any OAM message subject to OAM processing that does not contain an Authentication TLV or if the Authentication TLV does not verify.

9. Loopback Message

9.1. Loopback OAM Message format

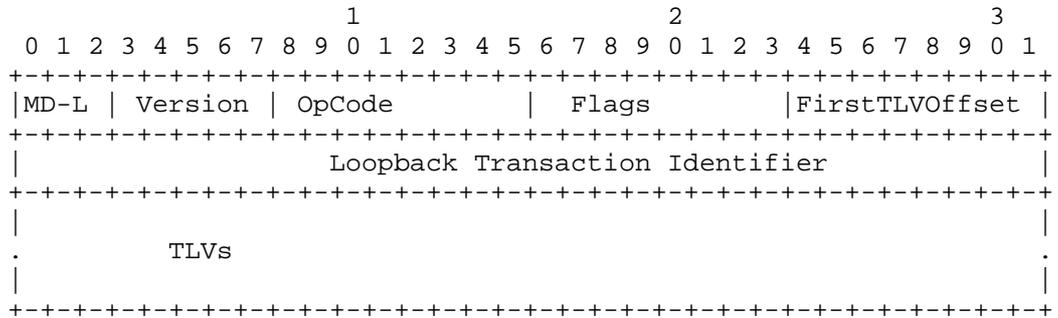


Figure 21 Loopback OAM Message Format

The above figure depicts the format of the Loopback Request and response messages as defined in [8021Q]. The Opcode for Loopback Message is set to 3 and the Opcode for the Reply Message is set to 2 [8021Q]. The Session Identification Number is a 32-bit

integer that allows the requesting RBridge to uniquely identify the corresponding session. Responding RBridges, without modification, MUST echo the received "Loopback Transaction Identifier" number.

9.2. Theory of Operation

9.2.1. Actions by Originator RBridge

The originator RBridge takes the following actions:

Identifies the destination RBridge nickname based on user specification or based on the specified destination MAC or IP address.

Constructs the flow entropy based on user specified parameters or implementation specific default parameters.

Constructs the TRILL OAM header: sets the opcode to Loopback message type (3)[8021Q]. Assigns applicable Loopback Transaction Identifier number for the request.

The TRILL OAM Application Identifier TLV MUST be included and with the flags set to applicable values.

Include following OAM TLVs, where applicable

- o Out of Band Reply Address TLV
- o Diagnostic Label TLV
- o Sender ID TLV

Specify the Hop count of the TRILL data frame per user specification or utilize an applicable Hop count value.

Dispatch the OAM frame for transmission.

RBridges may continue to retransmit the request at periodic intervals, until a response is received or the re-transmission count expires. At each transmission Session Identification number MUST be incremented.

9.2.2. Intermediate RBridge

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

9.2.3. Destination RBridge

If the Loopback message is addressed to the local RBridge and satisfies the OAM identification criteria specified in section 3.1. then, the RBridge data plane forwards the message to the CPU for further processing.

The TRILL OAM application layer further validates the received OAM frame by checking for the presence of OAM-EtherType at the end of the flow entropy. Frames that do not contain OAM-EtherType at the end of the flow entropy MUST be discarded.

Construction of the TRILL OAM response:

TRILL OAM application encodes the received TRILL header and flow entropy in the Original payload TLV and includes it in the OAM message.

Set the Return Code to (1) "Reply" and Return sub code to zero (0) "Valid Response". Update the TRILL OAM opcode to 2 (Loopback Message Reply)

Optionally, if the VLAN/FGL identifier value of the received flow entropy differs from the value specified in the diagnostic Label, set the Label Error Flag on TRILL OAM Application Identifier TLV.

Include the sender ID TLV (1)

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 IP forwarding process.

10. Path Trace Message

The primary use of the Path Trace Message is for fault isolation. It may also be used for plotting the path taken from a given RBridge to another RBridge.

[8021Q] accomplishes the objectives of the TRILL Path Trace Message using Link Trace Messages. Link Trace Messages utilize a well-known multicast MAC address. This works for [8021Q], because for 802.1 both the unicast and multicast paths are congruent. However, in TRILL multicast and unicast are not congruent. Hence, TRILL OAM uses a new message format: the Path Trace message.

The Path Trace Message has the same format as Loopback Message. The Opcode for Path Trace Reply is TBD1 and for Path Trace Message is TBD2.

Operation of the Path Trace message is identical to the Loopback message except that it is first transmitted with a TRILL Header Hop count field value of 1. The sending RBridge expects an "Intermediate RBridge" Return sub-code from the next hop or a "Valid response" Return sub-Code response from the destination RBridge. If an "Intermediate RBridge" Return sub-code is received in the response, the originator RBridge records the information received from intermediate node that generated the message and resends the message by incrementing the previous Hop count value by 1. This process is continued until, a response is received from the destination RBridge or Path Trace process timeout occur or Hop count reaches a configured maximum value.

10.1. Theory of Operation

10.1.1. Action by Originator RBridge

Identify the destination RBridge based on user specification or based on location of the specified MAC address.

Construct the flow entropy based on user specified parameters or implementation specific default parameters.

Construct the TRILL OAM header: Set the opcode to Path Trace Request message type (TBD2). Assign an applicable Session Identification number for the request. Return-code and sub-code MUST be set to zero.

The TRILL OAM Application Identifier TLV MUST be included and set the flags to applicable values.

Include following OAM TLVs, where applicable

- o Out of Band Reply Address TLV
- o Diagnostic Label TLV
- o Include the Sender ID TLV

Specify the Hop count of the TRILL data frame as 1 for the first request.

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

An RBridge may continue to retransmit the request at periodic intervals, until a response is received or the re-transmission count expires. At each new re-transmission, the Session Identification number MUST be incremented. Additionally, for responses received from intermediate RBridges, the RBridge nickname and interface information MUST be recorded.

10.1.2. Intermediate RBridge

Path Trace Messages transit through Intermediate RBridges transparently, unless Hop-count has expired.

TRILL OAM application layer further validates the received OAM frame by examining the presence of TRILL Alert Flag and OAM-Ethertype at the end of the flow entropy and by examining the MD Level. Frames that do not contain OAM-Ethertype at the end of the flow entropy MUST be discarded.

Construction of the TRILL OAM response:

TRILL OAM application encodes the received TRILL header and flow entropy in the Original payload TLV and include it in the OAM message.

Set the Return Code to (1) "Reply" and Return sub code to zero (2) "Intermediate RBridge". Update the TRILL OAM opcode to TBD1 (Path Trace Message Reply).

If the VLAN/FGL identifier value of the received flow entropy differs from the value specified in the diagnostic Label, set the Label Error Flag on TRILL OAM Application Identifier TLV.

Include following TLVs

Previous RBridge nickname TLV (69)

Reply Ingress TLV (5)

Reply Egress TLV (6)

Interface Status TLV (4)

TRILL Next Hop RBridge (Repeat for each ECMP) (70)

Sender ID TLV (1)

If Label error detected, set C flag (Label error detected) in the version.

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.

10.1.3. Destination RBridge

Processing is identical to section 10.1.2. With the exception that TRILL OAM Opcode is set to Path Trace Reply (TBD1).

11. Multi-Destination Tree Verification Message (MTVM)

Multi-Destination Tree Verification messages allow verifying TRILL distribution tree integrity and pruning. TRILL VLAN/FGL and multicast pruning are described in [RFC6325] [RFC7180] and [RFC7172]. Multi-destination tree verification and Multicast group verification messages are designed to detect pruning defects. Additionally, these tools can be used for plotting a given multicast tree within the TRILL campus.

Multi-Destination tree verification OAM frames are copied to the CPU of every intermediate RBridge that is part of the distribution tree being verified. The originator of the Multi-destination Tree verification message specifies the scope of RBridges from which a response is required. Only the RBridges listed in the scope field respond to the request. Other RBridges silently discard the request. Inclusion of the scope parameter is required to prevent receiving an excessive number of responses. The typical scenario of distribution tree verification or group verification, involves verifying multicast connectivity to a selected set of end-nodes as opposed to the entire network. Availability of the scope facilitates narrowing down the focus to only the RBridges of interest.

Implementations MAY choose to rate-limit CPU bound multicast traffic. As a result of rate-limiting or due to other congestion conditions, MTVM messages may be discarded from time to time by the intermediate RBridges and the requester may be required to retransmit the request. Implementations SHOULD narrow the

embedded scope of retransmission request only to RBridges that have failed to respond.

11.1. Multi-Destination Tree Verification Message (MTVM) Format

Format of MTVM is identical to that of Loopback Message format defined in section 9. with the exception that the Loopback Transaction Identifier, in section 9.1. , is replaced with the Session Identifier and the Op-Codes used is TBD4.

11.2. Theory of Operation

11.2.1. Actions by Originator RBridge

The user is required at a minimum to specify either the distribution trees that need to be verified, or the Multicast MAC address and VLAN/FGL, or VLAN/FGL and Multicast destination IP address. Alternatively, for more specific multicast flow verification, the user MAY specify more information e.g. source MAC address, VLAN/FGL, Destination and Source IP addresses. Implementations, at a minimum, must allow the user to specify a choice of distribution trees, Destination Multicast MAC address and VLAN/FGL that needs to be verified. Although, it is not mandatory, it is highly desired to provide an option to specify the scope. It should be noted that the source MAC address and some other parameters may not be specified if the Backwards Compatibility Method of Appendix A is used to identify the OAM frames.

Default parameters MUST be used for unspecified parameters. Flow entropy is constructed based on user specified parameters and/or default parameters.

Based on user specified parameters, the originating RBridge does the following:

- Identifies the nickname that represents the multicast tree.

- Obtains the applicable Hop count value for the selected multicast tree.

- Constructs TRILL OAM message header and include Session Identification number. Session Identification number facilitate the originator mapping the response to the correct request.

- Includes TRILL OAM Application Identifier TLV, which MUST be included.

Includes the Op-Code Multicast Tree Verification Message (TBD4)

Includes RBridge scope TLV (TBDe)

Optionally, include following TLV, where applicable

- o Out-of-band IP address (TBDb)
- o Diagnostic Label (TBDd)
- o Sender ID TLV (1)

Specify the Hop count of the TRILL data frame per user specification or alternatively utilize the applicable Hop count value if TRILL Hop count is not being specified by the user; and

Dispatch the OAM frame to the TRILL data plane to be ingressed for transmission.

The RBridge may continue to retransmit the request at a periodic interval until either a response is received or the re-transmission count expires. At each new re-transmission, the Session Identification number MUST be incremented. At each re-transmission, the RBridge may further reduce the scope to the RBridges that it has not received a response from.

11.2.2. Receiving RBridge

Receiving RBridges identify multicast verification frames per the procedure explained in sections 3.2.

The CPU of the RBridge validates the frame and analyzes the scope RBridge list. If the RBridge scope TLV is present and the local RBridge nickname is not specified in the scope list, it will silently discard the frame. If the local RBridge is specified in the scope list OR RBridge scope TLV is absent, the receiving RBridge proceeds with further processing as defined in section 11.2.3.

11.2.3. In scope RBridges

Construction of the TRILL OAM response:

TRILL OAM application encodes the received TRILL header and flow entropy in the Original payload TLV and includes them in the OAM message.

Set the Return Code to (0) and Return sub code to zero (0).
Update the TRILL OAM opcode to TBD3 (Multicast Tree Verification Reply).

Include following TLVs:

Previous RBridge nickname TLV (TBdf)

Reply Ingress TLV (5)

Interface Status TLV (4)

TRILL Next Hop RBridge List (TBDg)

Sender ID TLV (1)

Multicast Receiver Availability TLV (TBDh)

If a Label (VLAN or FGL) cross connect error is detected, set the C flag (Cross connect error detected) in the Application Identifier 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.

12. Application of Continuity Check Message (CCM) in TRILL

Section 7. provides an overview of CCM Messages defined in [8021Q] and how they can be used within the TRILL OAM. This section, presents the application and Theory of Operations of CCM within the TRILL OAM framework. Readers are referred to [8021Q] for CCM message format and applicable TLV definitions and usages. Only the TRILL specific aspects are explained below.

In TRILL, between any two given MEPs there can be multiple potential paths. Whereas in [8021Q], there is always a single path between any two MEPs at any given time. [RFC6905] requires solutions to have the ability to monitor continuity over one or more paths.

CCM Messages are uni-directional, such that there is no explicit response to a received CCM message. Connectivity status is

indicated by setting the applicable flags (e.g. RDI) of the CCM messages transmitted by an MEP.

It is important that the solution presented in this document accomplishes the requirements specified in [RFC6905] within the framework of [8021Q] in a straightforward manner and with minimum changes. Section 8 above defines multiple flows within the CCM object, each corresponding to a flow that a given MEP wishes to monitor. Hence, CCM, in multipath environments like TRILL, monitors per flow connectivity and cross connect errors.

Receiving MEPs do not cross check whether a received CCM belongs to a specific flow from the originating RBridge. Any attempt to track status of individual flows may explode the amount of state information that any given RBridge has to maintain.

The obvious question arises: How does the originating RBridge know which flow or flows are at fault?

This is accomplished with a combination of the RDI flag in the CCM header, flow-id TLV, and SNMP Notifications (Traps). Section 12.1. below discuss the procedure.

12.1.1. CCM Error Notification

Each MEP transmits 4 CCM messages per each flow. ([8021Q] detects CCM fault when 3 consecutive CCM messages are lost). Each CCM Message has a unique sequence number and unique flow-identifier. The flow identifier is included in the OAM message via flow-id TLV.

When an MEP notices a CCM timeout from a remote MEP (MEP-A), it sets the RDI flag on the next CCM message it generates. Additionally, it logs and sends SNMP notification that contain the remote MEP Identification, flow-id and the Sequence Number of the last CCM message it received and if available, the flow-id and the Sequence Number of the first CCM message it received after the failure. Each MEP maintains a unique flow-id per each flow, hence the operator can easily identify flows that correspond to the specific flow-id.

The following example illustrates the above.

Assume there are two MEPs, MEP-A and MEP-B.

Assume there are 3 flows between MEP-A and MEP-B.

Let's assume MEP-A allocates sequence numbers as follows

Flow-1 Sequence={1,2,3,4,13,14,15,16,.. } flow-id=(1)

Flow-2 Sequence={5,6,7,8,17,18,19,20,.. } flow-id=(2)

Flow-3 Sequence={9,10,12,11,21,22,23,24,.. } flow-id=(3)

Let's Assume Flow-2 is at fault.

MEP-B, receives CCM from MEP-A with sequence numbers 1,2,3,4, but did not receive 5,6,7,8. CCM timeout is set to 3 CCM intervals in [8021Q]. Hence MEP-B detects the error at the 8'th CCM message. At this time the sequence number of the last good CCM message MEP-B has received from MEP-A is 4 and flow-id of the last good CCM Message is (1). Hence MEP-B will generate a CCM error SNMP notification with MEP-A and Last good flow-id (1) and sequence number 4.

When MEP-A switches to flow-3 after transmitting flow-2, MEP-B will start receiving CCM messages. In the foregoing example it will be CCM message with Sequence Numbers 9,10,11,12,21 and so on. When in receipt of a new CCM message from a specific MEP, after a CCM timeout, the TRILL OAM will generate an SNMP Notification of CCM resume with remote MEP-ID and the first valid flow-id and the Sequence number after the CCM timeout. In the foregoing example, it is MEP-A, flow-id (3) and Sequence Number 9.

The remote MEP list under the CCM MIB Object is augmented to contain "Last Sequence Number", flow-id and "CCM Timeout" variables. Last Sequence Number and flow-id are updated every time a CCM is received from a remote MEP. CCM Timeout variable is set when the CCM timeout occurs and is cleared when a CCM is received.

12.2. Theory of Operation

12.2.1. Actions by Originator RBridge

Derive the flow entropy based on flow entropy specified in the CCM Management object.

Construct the TRILL CCM OAM header as specified in [8021Q].

TRILL OAM Version TLV MUST be included as the first TLV and set the flags to applicable values.

Include other TLVs specified in [8021Q]

Include the following optional TLV, where applicable

- o Sender ID TLV (1)

Specify the Hop count of the TRILL data frame per user specification or utilize an applicable Hop count value.

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

An RBridge transmits a total of 4 requests, each at CCM retransmission interval. At each transmission the Session Identification number MUST be incremented by one.

At the 5'th retransmission interval, flow entropy of the CCM packet is updated to the next flow entropy specified in the CCM Management Object. If current flow entropy is the last flow entropy specified, move to the first flow entropy specified and continue the process.

12.2.2. Intermediate RBridge

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

12.2.3. Destination RBridge

If the CCM Message is addressed to the local RBridge or multicast and satisfies OAM identification methods specified in sections 3.2. then the RBridge data plane forwards the message to the CPU for further processing.

The TRILL OAM application layer further validates the received OAM frame by examining the presence of OAM-Ethertype at the end of the flow entropy. Frames that do not contain OAM-Ethertype at the end of the flow entropy MUST be discarded.

Validate the MD-LEVEL and pass the packet to the Opcode de-multiplexer. The Opcode de-multiplexer delivers CCM packets to the CCM process.

The CCM Process performs processing specified in [8021Q].

Additionally the CCM process updates the CCM Management Object with the sequence number of the received CCM packet. Note: The last received CCM sequence number and CCM timeout are tracked per each remote MEP.

If the CCM timeout is true for the sending remote MEP, then clear the CCM timeout in the CCM Management object and generate the SNMP notification as specified above.

13. Fragmented Reply

TRILL OAM allows Fragmented reply messages. In case of Fragmented Replies, all part of the reply MUST follow the procedure defined in this section.

The same session Identification Number MUST be included in all related fragments of the same message.

The TRILL OAM Application Identifier TLV MUST be included, with fragment-ID field monotonically increasing with each fragment transmitted with the appropriate Final Flag field. The Final Flag, MUST, only be equal to one on the final fragment of the reply.

On the receiver, process MUST order the fragments based on the fragment id. Any fragments received after final fragment MUST be discarded. Messages with incomplete fragments (i.e. messages with one or missing fragments after the receipt of the fragment with the final flag set) MUST be discarded as well.

If number of fragments exceed the maximum supported fragments (255), then return code of MUST be set according to the message and return sub code MUST be set to 1 indicating fragment limit exceed.

14. Security Considerations

Forged OAM packets could cause false error or failure indications or mask actual errors or failures. For protection against forged OAM packets, the Authentication TLV (see Section 8.4.13) can be used in and OAM message in TRILL but is, of course, ineffective unless verified.

For general TRILL related security considerations, please refer to [RFC6325].

[8021Q] requires that the MEP filters or pass through OAM messages based on the MD-Level. The MD-Level is embedded deep in the OAM message. Hence, conventional methods of frame filtering may not be able to filter frames based on the MD-Level. As a result, OAM messages that must be dropped due to MD level mismatch may leak into a TRILL domain with different MD-Level.

This leaking may not cause any functionality loss. The receiving MEP/MIP is required to validate the MD-level prior to acting on the message. Any frames received with an incorrect MD-Level need to be dropped.

Generally, a single operator manages each TRILL campus, hence there is no risk of security exposure. However, in the event of multi operator deployments, operators should be aware of possible exposure of device specific information and appropriate measures must be taken.

It is also important to note that the MPLS OAM [RFC4379] framework does not include the concept of domains and OAM filtering based on operators. It is our opinion that the lack of OAM frame filtering based on domains does not introduce significant functional deficiency or security risk.

It is possible to mandate requiring different credentials to use different OAM functions or capabilities within a specific OAM function. Implementations may consider grouping users to different security clearance levels and restricting functions and capabilities to different clearance levels. However, Exact implementation details of such a framework are outside the scope of this document.

15. IANA Considerations

IANA is requested to assign the following:

15.1. OAM Capabilitiy Flags

Assign two TRILL-VER sub-TLV Capability Flags (see Section 3.3) as follows:

Bit	Description	Reference
---	-----	-----
TBD[2]	OAM capable	[this document]

TBD[3] Backwards compatible OAM [this document]

15.2. CFM Code Points

IANA is requested to assign four Op-Codes from the CFM OAM IETF Op-Codes sub-registry as follows [suggested values in square brackets]:

Value =====	Assignment =====	Reference =====
TBD1[64]	Path Trace Reply	[this document]
TBD2[65]	Path Trace Message	[this document]
TBD3[66]	Multicast Tree Verification Reply	[this document]
TBD4[67]	Multicast Tree Verification Messages	[this document]

IANA is requested to assign eleven TLV Types from the CFM OAM IETF TLV Types sub-registry as follows [suggested values in square brackets]:

Value =====	Assignment =====	Reference =====
TBDa[64]	TRILL OAM Application Identifier TLV	[this document]
TBDb[65]	Out of Band Reply Address TLV	[this document]
TBDc[66]	Diagnostic Label TLV	[this document]
TBDd[67]	Original Data Payload TLV	[this document]
TBDe[68]	RBridge Scope TLV	[this document]
TBDf[69]	Previous RBridge nickname TLV	[this document]
TBDg[70]	Next Hop RBridge List TLV	[this document]
TBDh[71]	Multicast Receiver Port count TLV	[this document]
TBDi[72]	Flow Identifier TLV	[this document]
TBDj[73]	Reflector Entropy TLV	[this document]
TBDk[74]	Authentication TLV	[this document]

15.3. MAC Addresses

IANA is requested to assigned a unicast and a multicast MAC address under the IANA OUI, for identification of OAM packets as discussed for the backward compatibility method (Appendix A, Section A.2) based on the request template in Appendix C. The assigned addresses are TBDmac1 and TBDmac2.

15.4. Return codes and sub codes

IANA is requested to create TRILL OAM Return code registry and, within each return code a separate sub code registry. Within TRILL OAM Return code and sub code registry following values are requested to be allocated.

Return code zero (0) is reserved for request messages.

Return code one (1) indicate reply message. Sub code zero (0) indicates valid response, sub code of one (1) indicates fragment limit exceeded Section 13.

All other Return code values from 2 to 255 are currently reserved and may be assigned in the future.

Return Code	Assignment	Sub codes
=====	=====	=====
0	Request message	0 - Valid request 1-255 Reserved
1	Reply message	0 - Valid response 1 - Fragment limit exceed 2 - Intermediate RBridge 3-255 Reserved

16. References

16.1. Normative References

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

- [RFC5310] Bhatia, M., "IS-IS Cryptographic Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC6325] Perlman, R., et.al., "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.
- [RFC7172] Eastlake 3rd, D., Zhang, M., Agarwal, P., Perlman, R., and D. Dutt, "Transparent Interconnection of Lots of Links (TRILL): Fine-Grained Labeling", RFC 7172, May 2014..
- [8021Q] IEEE, "Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August, 2011.
- [IS-IS] ISO/IEC 10589:2002, Second Edition, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2002.

16.2. Informative References

- [RFC4379] Kompella, K. et.al, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC6291] Andersson, L., et.al., "Guidelines for the use of the "OAM" Acronym in the IETF" RFC 6291, June 2011.
- [RFC6361] Carlson, J. and Eastlake, D. "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011.
- [RFC6905] Senevirathne, T. et.al "Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)", RFC 6905, March 2013.
- [RFC7176] Eastlake, D., Senevirathne, T., Ghanwani, A., Dutt, D., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", RFC 7176 May 2014.
- [RFC7180] Eastlake, Donald, et.al. "TRILL: Clarifications, Corrections, and Updates", RFC 7180 May 2014.

- [RFC7174] Salam, S., et.al., "TRILL OAM Framework", RFC 7174 , May 2014.
- [RFC7179] Eastlake, Donald, et.al. "TRILL: Header Extension", RFC 7179, May 2014.
- [Y1731] ITU, "OAM functions and mechanisms for Ethernet based networks", ITU-T G.8013/Y.1731, July, 2011.
- [RFC7178] D. Eastlake, et.al. , "TRILL: RBridge Channel Support", RFC 7178, May 2014.
- [TRILLOAMMIB] Deepak Kumar et.al, "TRILL OAM MIB", draft-deepak-trill-oam-mib, May 2013, work in progress.

17. Acknowledgments

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

Acknowledgments are due for many who volunteered to review this document, notably, Dan Romascanu, Gayle Nobel and Tal Mizrahi.

Special appreciations are due for Dinesh Dutt for his support and encouragement, especially during the initial discussion phase of TRILL OAM.

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

Appendix A.

Backwards Compatibility

Methodology presented above in this document is in-line with the [8021Q] framework for providing fault management coverage. However, in practice, some TRILL platforms may not have the capabilities to support some of the required techniques. In this section, we present a method that allows R Bridges, which do not have the required hardware capabilities, to participate in the TRILL OAM solution.

There are two broad areas to be considered; 1. Maintenance Point (MEP/MIP) Model 2. Data plane encoding and frame identification

A.1 Maintenance Point (MEP/MIP) Model

For backwards compatibility, MEPs and MIPs are located in the CPU. This will be referred to as the "central brain" model as opposed to "port brain" model.

In the "central brain" model, an R Bridge using either ACLs or some other method, forwards qualifying OAM messages to the CPU. The CPU then performs the required processing and multiplexing to the correct MP (Maintenance Point).

Additionally, R Bridges MUST have the capability to prevent the leaking of OAM packets, as specified in [RFC6905].

A.2 Data plane encoding and frame identification

The backwards compatibility method presented in this section defines methods to identify OAM frames when implementations do not have capabilities to utilize TRILL OAM Alert flag presented earlier to identify OAM frames, in the hardware.

It is assumed ECMP path selection of non-IP flows utilize MAC DA, MAC SA and VLAN, IP Flows utilize IP DA, IP SA and TCP/UDP port numbers and other Layer 3 and Layer 4 information. The well-known fields to identify OAM flows are chosen such that they mimic the ECMP selection of the actual data along the path. However, it is important to note that, there may be implementations that would utilize these well-known fields for ECMP selections. Hence, implementations that support OAM SHOULD move to utilizing TRILL Alert Flag, as soon as possible and methods presented here SHOULD be used only as an interim solution.

Identification methods are divided in to 4 broader groups:

1. Identification of Unicast non-IP OAM Flows,
2. Identification of Multicast non-IP OAM Flows,
3. Identification of Unicast IP OAM Flows and
4. Identification of Multicast IP OAM Flows

As presented in the table below, based on the flow type (as defined above), implementations are required to use a well-known value in either the Inner.MacSA field or OAM Ethertype field to identify OAM flows.

Receiving RBridge identifies OAM flows based on the presence of the well-known values in the specified fields, and additionally, for unicast flows, egress RBridge nickname of the packet MUST match that of the local RBridge or for multicast flows, TRILL header mutlicast flag MUST be set.

Unicast OAM flows that qualify for local processing MUST be redirected to the OAM process and MUST NOT be forwarded (that to prevent leaking of the packet out of the TRILL campus).

A copy of Multicast OAM flows that qualify for local processing MUST be sent to the OAM process and packet MUST be forwarded along the normal path. Additionally, methods MUST be in place to prevent multicast packets leaking out of the TRILL campus.

The following table summarizes the identification of different OAM frames from data frames.

Flow Entropy	Inner MacSA	OAM Ether Type	Egress nickname
unicast no IP	N/A	Match	Match
Multicast no IP	N/A	Match	N/A
Unicast IP	Match	N/A	Match
Multicast IP	Match	N/A	N/A

Figure 22 Identification of TRILL OAM Frames

The unicast and multicast Inner.MacSAs used for the unicast and multicast IP cases, respectively, are TBDmac1 and TBDmac2 assigned by the request in Appendix C.

It is important to note that all RBridges MUST generate OAM flows with "A" flag set and CFM EtherType "0x8902" at the flow entropy off-set. However, well-known values MUST be utilized as part of the flow-entropy when generating OAM messages destined for older RBridges that are compliant to the backwards compatibility method defined in this appendix.

Appendix B.

Base Mode for TRILL OAM

CFM, as defined in [8021Q], requires configuration of several parameters before the protocol can be used. These parameters include MAID, Maintenance Domain Level (MD-LEVEL) and MEPIDs. The Base Mode for TRILL OAM defined here facilitates ease of use and provides out of the box plug-and-play capabilities, supporting the Operational and Manageability considerations described in Section 6 of [RFC7174].

All Rbridges that support TRILL OAM MUST support Base Mode operation.

All Rbridges MUST create a default MA with MAID as specified herein.

MAID [8021Q] has a flexible format and includes two parts: Maintenance Domain Name and Short MA name. In the Based Mode of operation, the value of the Maintenance Domain Name must be the character string "TrillBaseMode" (excluding the quotes "). In Base Mode operation Short MA Name format is set to 2-octet integer format (value 3 in Short MA Format field) and Short MA name set to 65532 (0xFFFC).

The Default MA belongs to MD-LEVEL 3.

In the Base Mode of operation, each RBridge creates a single UP MEP associated with a virtual OAM port with no physical layer (NULL PHY). The MEPID associated with this MEP is the 2-octet RBridge Nickname.

By default, all Rbridges operating in the Base Mode for TRILL OAM are able to initiate LBM, PT and other OAM tools with no configuration.

Implementations MAY provide default flow-entropy to be included in OAM messages. Content of the default flow-entropy is outside the scope of this document.

Figure 23, below depicts encoding of MAID within CCM messages.

```

+-----+-----+
|Field Name   |Size   |
+-----+-----+
|Maintenance  |  1    |
|Domain Format |       |
+-----+-----+
|Maintenance  |  2    |
|Domain Length|       |
+-----+-----+
|Maintenance  |variable|
|Domain Name  |       |
+-----+-----+
|Short MA     |  1    |
|Name Format   |       |
+-----+-----+
|Short MA     |  2    |
|Name Length  |       |
+-----+-----+
|Short MA     |variable|
|Name        |       |
+-----+-----+
|Padding      |Variable|
+-----+-----+

```

Figure 23 MAID structure as defined in [8021Q]

Maintenance Domain Name Format is set to Value: 4
 Maintenance Domain Name Length is set to value: 13
 Maintenance Domain Name is set to: TrillBaseMode
 Short MA Name Format is set to value: 3
 Short MA Name Length is set to value: 2
 Short MA Name is set to: FFFC
 Padding: set of zero up to 48 octets of total length of the MAID.
 Please refer to [8021Q] for details.

Appendix C.

MAC Addresses Request

Applicant Name: IETF TRILL Working Group

Applicant Email: tsenevir@cisco.com

Applicant Telephone: +1-408-853-2291

Use Name: TRILL OAM

Document: draft-tissa-trill-oam-fm

Specify whether this is an application for EUI-48 or EUI-64

identifiers: EUI-48

Size of Block requested: 1

Specify multicast, unicast, or both: Both

Authors' Addresses

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

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

Norman Finn
CISCO Systems
510 McCarthy Blvd
Milpitas, CA 95035
USA

Email: nfinn@cisco.com

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

Email: ssalam@cisco.com

Deepak Kumar
CISCO Systems
510 McCarthy Blvd,
Milpitas, CA 95035, USA

Phone : +1 408-853-9760
Email: dekumar@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 95951
USA

Email: aldrin.ietf@gmail.com

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China

Phone: +86-25-56625375
Email: liyizhou@huawei.com

