

Internet Engineering Task Force (IETF)
Request for Comments: 7455
Updates: 6325
Category: Standards Track
ISSN: 2070-1721

T. Senevirathne
N. Finn
S. Salam
D. Kumar
Cisco
D. Eastlake 3rd
S. Aldrin
Y. Li
Huawei
March 2015

Transparent Interconnection of Lots of Links (TRILL): Fault Management

Abstract

This document specifies Transparent Interconnection of Lots of Links (TRILL) Operations, Administration, and Maintenance (OAM) fault management. Methods in this document follow the CFM (Connectivity Fault Management) framework defined in IEEE 802.1 and reuse OAM tools where possible. Additional messages and TLVs are defined for TRILL-specific applications or for cases where a different set of information is required other than CFM as defined in IEEE 802.1. This document updates RFC 6325.

Status of This Memo

This is an Internet Standards Track document.

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

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7455>.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	5
2. Conventions Used in This Document	5
3. General Format of TRILL OAM Packets	6
3.1. Identification of TRILL OAM Frames	8
3.2. Use of TRILL OAM Alert Flag	8
3.2.1. Handling of TRILL Frames with the "A" Flag	9
3.3. OAM Capability Announcement	9
3.4. Identification of the OAM Message	10
4. TRILL OAM Layering vs. IEEE Layering	11
4.1. Processing at the ISS Layer	12
4.1.1. Receive Processing	12
4.1.2. Transmit Processing	12
4.2. End-Station VLAN and Priority Processing	12
4.2.1. Receive Processing	12
4.2.2. Transmit Processing	12
4.3. TRILL Encapsulation and Decapsulation Layer	12
4.3.1. Receive Processing for Unicast Packets	12
4.3.2. Transmit Processing for Unicast Packets	13
4.3.3. Receive Processing for Multicast Packets	14
4.3.4. Transmit Processing of Multicast Packets	15
4.4. TRILL OAM Layer Processing	16
5. Maintenance Associations (MAs) in TRILL	17
6. MEP Addressing	18
6.1. Use of MIP in TRILL	21
7. Continuity Check Message (CCM)	22
8. TRILL OAM Message Channel	25
8.1. TRILL OAM Message Header	25
8.2. TRILL-Specific OAM OpCodes	26
8.3. Format of TRILL OAM TLV	26
8.4. TRILL OAM TLVs	27
8.4.1. Common TLVs between CFM and TRILL	27
8.4.2. TRILL OAM-Specific TLVs	27
8.4.3. TRILL OAM Application Identifier TLV	28
8.4.4. Out-of-Band Reply Address TLV	30
8.4.5. Diagnostic Label TLV	31
8.4.6. Original Data Payload TLV	32
8.4.7. RBridge Scope TLV	32
8.4.8. Previous RBridge Nickname TLV	33
8.4.9. Next-Hop RBridge List TLV	34
8.4.10. Multicast Receiver Port Count TLV	34
8.4.11. Flow Identifier TLV	35
8.4.12. Reflector Entropy TLV	36
8.4.13. Authentication TLV	37

9. Loopback Message	38
9.1. Loopback Message Format	38
9.2. Theory of Operation	39
9.2.1. Actions by Originator RBridge	39
9.2.2. Intermediate RBridge	39
9.2.3. Destination RBridge	40
10. Path Trace Message	40
10.1. Theory of Operation	41
10.1.1. Actions by Originator RBridge	41
10.1.2. Intermediate RBridge	42
10.1.3. Destination RBridge	43
11. Multi-Destination Tree Verification Message (MTVM)	43
11.1. MTVM Format	44
11.2. Theory of Operation	44
11.2.1. Actions by Originator RBridge	44
11.2.2. Receiving RBridge	45
11.2.3. In-Scope RBridges	45
12. Application of Continuity Check Message (CCM) in TRILL	46
12.1. CCM Error Notification	47
12.2. Theory of Operation	48
12.2.1. Actions by Originator RBridge	48
12.2.2. Intermediate RBridge	49
12.2.3. Destination RBridge	49
13. Fragmented Reply	50
14. Security Considerations	50
15. IANA Considerations	52
15.1. OAM Capability Flags	52
15.2. CFM Code Points	52
15.3. MAC Addresses	53
15.4. Return Codes and Sub-codes	53
15.5. TRILL Nickname Address Family	54
16. References	54
16.1. Normative References	54
16.2. Informative References	55
Appendix A. Backwards Compatibility	57
A.1. Maintenance Point (MEP/MIP) Model	57
A.2. Data-Plane Encoding and Frame Identification	57
Appendix B. Base Mode for TRILL OAM	59
Appendix C. MAC Addresses Request	61
Acknowledgments	62
Authors' Addresses	62

1. Introduction

The general structure of TRILL OAM messages is presented in [RFC7174]. TRILL OAM messages consist of six parts: Link Header, TRILL Header, Flow Entropy, OAM Ethertype, 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 that readers are familiar with [8021Q] and [Y1731]. Readers who are not familiar with these documents are encouraged to review them.

This document specifies TRILL OAM fault management. It updates [RFC6325] as specified in Section 3.1. TRILL performance monitoring is specified in [RFC7456].

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].

Capitalized IANA Considerations terms such as "Standards Action" are to be interpreted as described in [RFC5226].

Acronyms used in the document include the following:

- CCM - Continuity Check Message [8021Q]
- DA - Destination Address
- ECMP - Equal-Cost Multipath
- FGL - Fine-Grained Label

ISS - Internal Sub-Layer Service [8021Q]
LBM - Loopback Message [8021Q]
LBR - Loopback Reply [8021Q]
MA - Maintenance Association [8021Q] [RFC7174]
MAC - Media Access Control (MAC)
MD - Maintenance Domain [8021Q]
MEP - Maintenance End Point [RFC7174] [8021Q]
MIP - Maintenance Intermediate Point [RFC7174] [8021Q]
MP - Maintenance Point [RFC7174]
MTVM - Multi-destination Tree Verification Message
MTVR - Multi-destination Tree Verification Reply
OAM - Operations, Administration, and Maintenance [RFC6291]
PRI - Priority of Ethernet Frames [8021Q]
PTM - Path Trace Message
PTR - Path Trace Reply
SA - Source Address
SAP - Service Access Point [8021Q]
TRILL - Transparent Interconnection of Lots of Links [RFC6325]

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 OAM messages. The details of the TRILL OAM frame format are defined in this document.

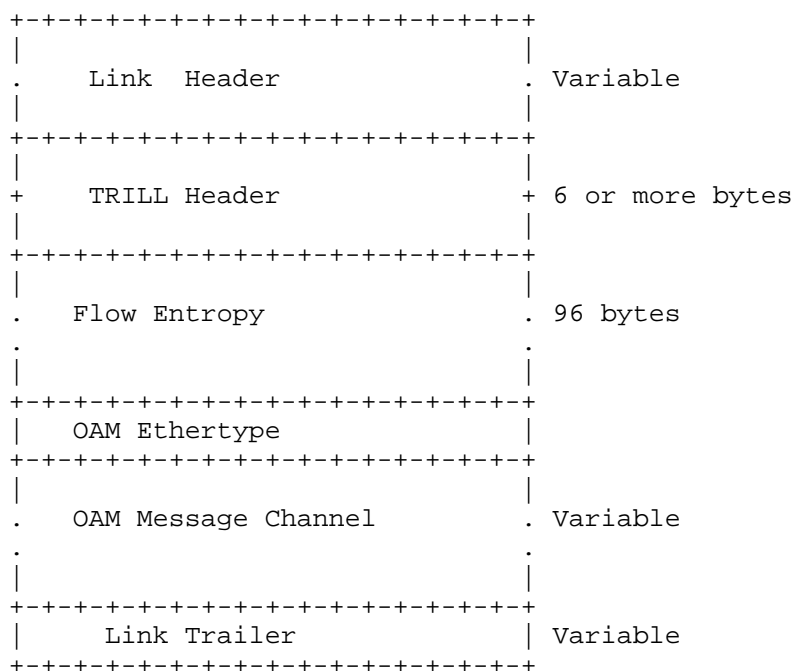


Figure 1: Format of TRILL OAM Messages

- o Link Header: Media-dependent header. For Ethernet, this includes the Destination MAC, Source MAC, VLAN (optional), and Ethertype fields.
- o TRILL Header: Fixed size of 6 bytes when the Extended Header is not included [RFC6325].
- o Flow Entropy: 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 information 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.

- o OAM Ethertype: A 16-bit Ethertype that identifies the OAM Message Channel that follows. This document specifies using the Ethertype 0x8902 allocated for CFM [8021Q].
- o OAM Message Channel: A variable-size section that carries OAM-related information. The message format is as specified in [8021Q].
- o 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 method to identify OAM frames. This document updates [RFC6325] to include specific methods to identify TRILL OAM frames. Section 3.2 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 the Version field in the TRILL Header as the Alert flag. The 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 the "A" flag and CFM Ethertype to identify TRILL OAM frames.

```

+-----+
| V | A | R | M | Op-Length | Hop Count |
+-----+
| Egress RBridge Nickname | Ingress RBridge Nickname |
+-----+
| Options...
+-----+

```

Figure 2: TRILL Header with the "A" Flag

- o 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

The 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 Ethertype (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 a CFM Ethertype are not considered OAM frames. Such frames MUST be silently discarded.

OAM-capable RBridges MUST NOT generate OAM frames to an RBridge that is not OAM capable.

Intermediate RBridges that are not OAM capable (i.e., do not understand the "A" flag) follow the process defined in Section 3.3 of [RFC6325] and forward OAM frames with the "A" flag unaltered.

3.3. OAM Capability Announcement

Any given RBridge can be (1) OAM incapable, (2) OAM capable with new extensions, or (3) OAM capable with the backwards-compatibility 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.

The capability flags defined in the 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 the "O" flag set to 1 and the "B" flag set to 1 indicates that the originating RBridge is OAM capable but utilizes the backwards-compatibility method defined in Appendix A. A capability announcement with the "O" flag set to 1 and the "B" flag set to 0 indicates that the originating RBridge is OAM capable and utilizes the method specified in Section 3.2.

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

```

+-----+
| Type           | (1 byte)
+-----+
| Length         | (1 byte)
+-----+
| Max-version    | (1 byte)
+-----+
|A|F|O|B|Other Capabilities and Header Flags| (4 bytes)
+-----+
0           1           3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 0 1

```

Figure 3: TRILL-VER Sub-TLV [RFC7176] with "O" and "B" Flags

In Figure 3, "A" is the Affinity sub-TLV support flag as indicated in [RFC7176], and "F" is the FGL-safe flag as indicated in [RFC7172] and [RFC7176]. The "O" and "B" flags are located after the "F" flag in the Capability and Header Flags field of the TRILL-VER sub-TLV, as depicted in Figure 3 above. Usage of the "O" and "B" flags is discussed above.

Absence of the TRILL-VER sub-TLV means the announcing RBridge is not OAM capable.

3.4. Identification of the OAM Message

The ingress RBridge nickname allows recipients to identify the origin of the message in most cases. However, when an out-of-band reply is generated, the responding RBridge nickname is not easy to identify.

The [8021Q] Sender ID TLV (1) provides methods to identify the device by including the Chassis ID. The Chassis ID allows different addressing formats such as IANA Address Family enumerations. IANA has allocated Address Family Number 16396 for TRILL nickname. In TRILL OAM, the Chassis ID sub-type of the Sender ID TLV is set to 16396, and the Chassis ID field contains the corresponding TRILL nickname.

When the Sender ID TLV is present and the Chassis ID sub-type is set to 16396, the sender RBridge TRILL nickname SHOULD be derived from the nickname embedded in the Chassis ID. Otherwise, the sender RBridge TRILL nickname SHOULD be derived from the ingress RBridge nickname.

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.

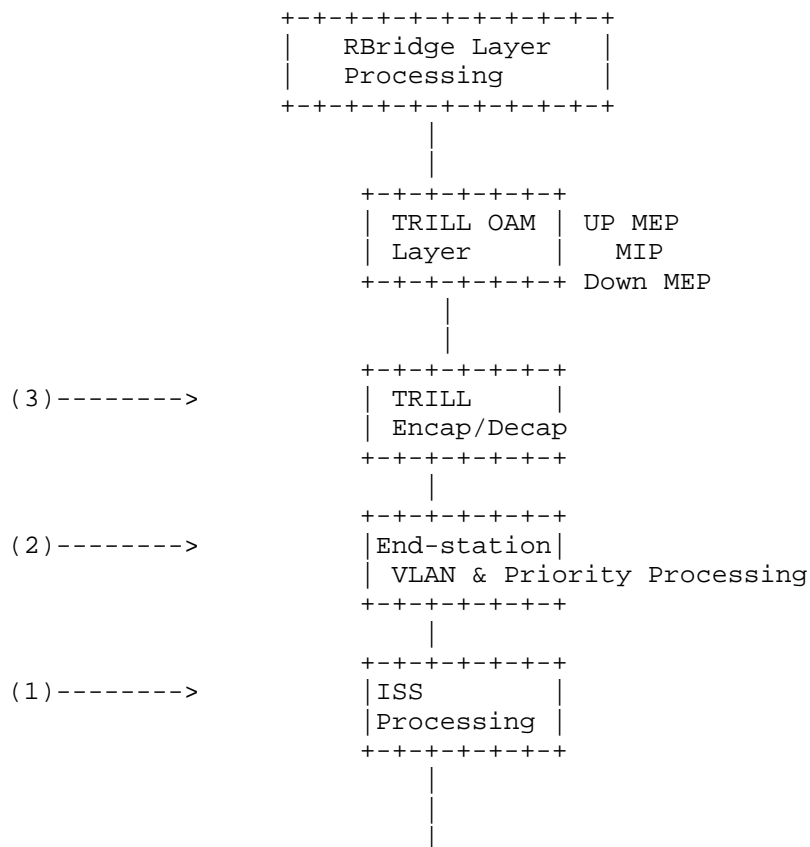


Figure 4: Placement of TRILL MP within IEEE 802.1

[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 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 the ISS Layer

4.1.1. Receive Processing

The ISS layer receives an indication from the port. It extracts DA and SA, and it marks the remainder of the payload as M1. The 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

Receive (DA, SA, M1) indication from the ISS layer. Extract 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 construct (DA, SA, VLAN, PRI, M2). VLAN+PRI+M2 maps to M1 in the received indication. Pass (DA, SA, VLAN, PRI, M2) to the TRILL Encapsulation/Decapsulation layer.

4.2.2. Transmit Processing

Receive (DA, SA, VLAN, PRI, M2) indication from the TRILL Encapsulation/Decapsulation layer. Merge VLAN, PRI, M2 to form M1. Pass down (DA, SA, M1) to the ISS layer.

4.3. TRILL Encapsulation and Decapsulation Layer

4.3.1. Receive Processing for Unicast Packets

- o Receive indication (DA, SA, VLAN, PRI, M2) from the End-Station VLAN and Priority Processing layer.
- o If the DA matches the port Local DA and the frame is of TRILL Ethertype:

- Discard DA, SA, VLAN, and PRI. From M2, derive (TRILL-HDR, iDA, iSA, i-VL, M3).
- If TRILL nickname is Local and TRILL Header Alert flag is set:
 - * Pass on to OAM processing.
- Else, pass on (TRILL-HDR, iDA, iSA, i-VL, M3) to the RBridge layer.
- o If the DA matches the port Local DA and the Ethertype is RBridge-Channel [RFC7178]:
 - Process as a possible unicast native RBridge Channel packet.
- o If the DA matches the port Local DA and the Ethertype is neither TRILL nor RBridge-Channel:
 - Discard packet.
- o If the DA does not match, the port is Appointed Forwarder for VLAN, and the Ethertype is not TRILL or RBridge-Channel:
 - Insert TRILL-HDR and send (TRILL-HDR, iDA, iSA, i-VL, M3) indication to the 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 the RBridge layer.
- o If the egress TRILL nickname is local:
 - If the port is Appointed Forwarder for iVL, the port is not configured as a trunk or point-to-point (P2P) port, the TRILL Alert flag is set, and the OAM Ethertype is present, then:
 - * Strip TRILL-HDR and construct (DA, SA, VLAN, M2) (this is the TRILL Egress Function).
 - Else:
 - * Discard packet.

- o If the egress TRILL nickname is not local:
 - Insert Outer.MacDA, Outer.MacSA, Outer.VLAN, and 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 End-Station VLAN and Priority Processing layer.

4.3.3. Receive Processing for Multicast Packets

- o Receive (DA, SA, V, M2) from the End-Station VLAN and Priority Processing layer.
- o If the DA is All-RBridges and the Ethertype is TRILL:
 - Strip DA, SA, and V. From M2, extract (TRILL-HDR, iDA, iSA, iVL, and M3).
 - If the TRILL Alert flag is set and the OAM Ethertype is present at the end of Flow Entropy:
 - * Perform OAM processing.
 - Else, extract the TRILL Header, inner MAC addresses, and Inner.VLAN, and pass indication (TRILL-HDR, iDA, iSA, iVL and M3) to the 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 the Ethertype is not TRILL or L2-IS-IS respectively:
 - 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:
 - Discard the packet.
- o If the DA is All-Edge-RBridges and the Ethertype is RBridge-Channel [RFC7178]:
 - 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:
 - 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 include the Link Aggregation Protocol and the Link-Layer Discovery Protocol).
- o If the DA is some other multicast value:
 - Insert TRILL-HDR and construct (TRILL-HDR, iDA, iSA, iVL, M3).
 - Pass the (TRILL-HDR, iDA, iSA, iVL, M3) to the 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 the RBridge layer.
- o If the TRILL Header multicast ("M") flag is set, the TRILL-HDR Alert flag is set, and the OAM Ethertype is present, then:
 - Construct (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: A second copy of native format is not made.
- o Else, if the TRILL Header multicast ("M") flag is set and the Alert flag not set:
 - 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, and iVL and construct (DA, SA, V, M2) for native format.
 - 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 the End-Station VLAN and Priority Processing layer.

4.4. TRILL OAM Layer Processing

The TRILL OAM layer is located between the TRILL Encapsulation/Decapsulation layer and the RBridge layer. It performs the following: 1) identifies OAM frames that need local processing and 2) performs OAM processing or redirects to the CPU for OAM processing.

- o Receive indication (TRILL-HDR, iDA, iSA, iVL, M3) from the RBridge layer. M3 is the payload after Inner.VLAN iVL.
- o If the TRILL Header multicast ("M") flag is set, the TRILL Alert flag is set, and TRILL OAM Ethertype is present, then:
 - 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:
 - + 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.
 - Else, if the TRILL Alert flag is set and the TRILL OAM Ethertype is present, then:
 - If MEP or MIP is configured on the Inner.VLAN/FGL of the packet, then:
 - * Discard packets that have MD-Level not present or where MD-Level is less than that of the MEP.
 - * If MD-Level matches MD-Level of the MEP, then:
 - + 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, for a non-OAM packet:
 - Continue.
- o Pass the indication (DA, SA, V, M2) to the End-Station VLAN and Priority Processing layer.

Note: In the receive path, the processing above compares with the Down MEP and MIP Half functions. In the transmit processing, it compares with Up MEP and MIP Half functions.

Appointed Forwarder is a function that the TRILL Encapsulation/Decapsulation layer performs. The TRILL Encapsulation/Decapsulation layer is responsible for prevention of leaking of OAM packets as native frames.

5. Maintenance Associations (MAs) 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 depicts the augmentation of the CFM MIB to add the TRILL-specific Flow Entropy.

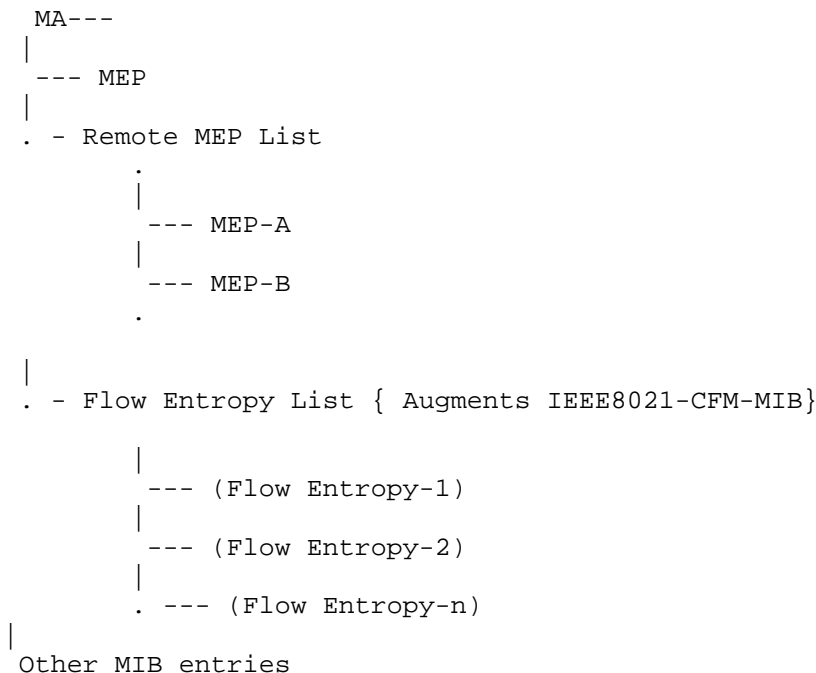


Figure 5: Correlation of TRILL-Augmented MIB

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 a combination of the egress RBridge nickname and the Inner.VLAN/FGL.

Additionally, MEPs are represented by a 2-octet MEP-ID that is independent of the underlying technology. In CFM [8021Q], the value of MEP-ID is restricted to the range of 1 to 8191. However, on a CFM [8021Q] packet, MEP-IDs are encoded as a 2-octet field. In the TRILL Base Mode operation presented in Appendix B, MEP-IDs are mapped 1-to-1 with the RBridge nicknames. Hence, in TRILL, a MEP-ID MUST be a number in the range from 1 to 65535.

At the MEP, OAM packets go through a hierarchy of OpCode demultiplexers. The OpCode demultiplexers channel the incoming OAM packets to the appropriate message processor (e.g., LBM). Refer to Figure 6 for a visual depiction of these different demultiplexers.

The demultiplexing sequence is as follows:

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. Redirects to the MD-Level demultiplexer.
3. The MD-Level demultiplexer 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 they belong to different MAs.)
 - a. If the packet MD-Level is equal to the configured MD-Level of the MEP, then pass to the OpCode demultiplexer.
 - b. If the packet MD-Level is less than the configured MD-Level of the MEP, discard the packet.
 - c. If the packet MD-Level is greater than the configured MD-Level of the MEP, then pass on to the next-higher MD-Level demultiplexer, if available. Otherwise, if no such higher MD-Level demultiplexer exists, then forward the packet as normal data.
4. The OpCode demultiplexer compares the OpCode in the packet with supported OpCodes.
 - a. If the OpCode is CCM, LBM, LBR, PTM, PTR, MTVM, or MTRV, then pass on to the correct processor.
 - b. If the OpCode is unknown, then discard.

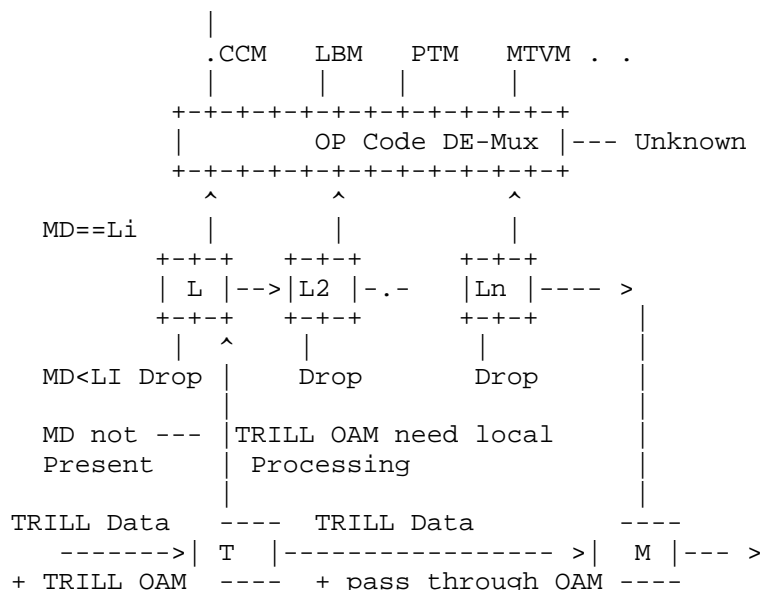


Figure 6: OAM Demultiplexers at MEP for Active SAP

- o T: Denotes Tap. Identifies OAM frames that need local processing. These are the packets with the Alert flag set and OAM Ethertype present after the Flow Entropy of the packet.
- o M: The post-processing merge that 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.
- o L: Denotes MD-Level processing. Packets whose MD-Level is less than the MD-Level of the current processing step will be dropped. Packets with equal MD-Levels are passed on to the OpCode demultiplexer. 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 demultiplexers. 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 (MIPs) 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. Any MIP residing within the ingress interface may reply to the incoming Path Trace Message without loss of functionality or information. As specified in Section 3.4, the address of the responding RBridge can be identified by means of the Sender ID TLV (1). The Reply Ingress TLV (5) identifies the interface id. The combination of these allows the recipient of the response to uniquely identify the responder.

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 Demultiplexers at MIP for Active SAP

- o T: Tap processing for MIP. All packets with the TRILL Header Alert flag set are captured.
- o L: MD-Level Processing. Packets with matching MD-Levels are "copied" to the OpCode demultiplexer, and the 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 OpCode demultiplexer.
- o M: The intermediate point processing merge that merges data and OAM messages that are passed through.

Packets with unknown OpCodes are counted and discarded.

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 MEP-ID 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 (MA) identifies the interested Maintenance 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 for the theory of operation of CCM.

The MIB in [8021Q] is augmented with the definition of Flow Entropy. Please see [TRILLOAMMIB] for this and other TRILL-related OAM MIB definitions. Figure 8 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 8021ag entries)
    |
    | --- (Hop Count) { Augments IEEE8021-CFM-MIB}
    |
    | --- (Any other 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 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 reusing the OAM message header defined in [8021Q].




Figure 9: OAM Message Format

- o MD-L: Maintenance Domain Level (3 bits). For TRILL, in general, this field is set to a single value across the TRILL campus. When using the 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 the 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 OpCode: Operation Code (8 bits). Specifies the operation performed by the message. See Section 8.2.
- 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.)
- o OpCode-Specific Information: May contain Session Identification Number, timestamp, etc.

The MD-L, Version, OpCode, Flags, and FirstTLVOffset fields collectively are referred to as the OAM message header.

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 Sections 10 and 11.

TRILL OAM message OpCodes:

- 64: Path Trace Reply
- 65: Path Trace Message
- 66: Multi-destination Tree Verification Reply
- 67: Multi-destination 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

- o Type (1 octet): Specifies the type of the TLV (see Section 8.4 for TLV types).
- o Length (2 octets): Specifies the length of the Value field in octets. Length of the Value field can be zero or more octets.

- o Value (variable): The length and the content of this field depend on the type of TLV. Please refer to applicable TLV definitions for 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. TLVS defined in [8021Q] are reused, where applicable.

8.4.1. Common TLVs between CFM and TRILL

The following TLVs are defined in [8021Q]. We reuse 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
64	TRILL OAM Application Identifier TLV
65	Out-of-Band Reply Address TLV
66	Diagnostic Label TLV
67	Original Data Payload TLV
68	RBridge Scope TLV
69	Previous RBridge Nickname TLV
70	Next-Hop RBridge List TLV
71	Multicast Receiver Port Count TLV
72	Flow Identifier TLV
73	Reflector Entropy TLV
74	Authentication TLV

The TRILL OAM Application Identifier TLV (64) MUST be the first TLV. An End TLV (0) MUST be included as the last TLV. All other TLVs can be included in any order.

8.4.3. TRILL OAM Application Identifier TLV

The TRILL OAM Application Identifier TLV carries information specific to TRILL OAM applications. 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 a TRILL MP.

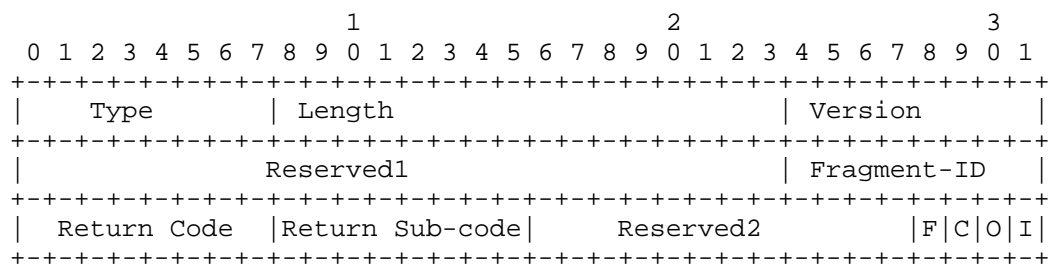


Figure 11: TRILL OAM Application Identifier TLV

- o Type (1 octet): 64, TRILL OAM Application Identifier TLV
- o Length (2 octets): 9
- o Version (1 octet): Currently set to zero. Indicates the TRILL OAM version. The TRILL OAM version can be different than the [8021Q] version.
- o Reserved1 (3 octets): Set to zero on transmission and ignored on reception.

- o Fragment-ID (1 octet): Indicates the fragment number of the current message. This applies only to reply messages; in request messages, it must be set to zero on transmission and ignored on receipt. The "F" flag defined below MUST be set with the final message, whether it is the last fragment of the fragmented message or the only message of the reply. Section 13 provides more details on OAM message fragmentation.
- o Return Code (1 octet): Set to zero on requests. Set to an appropriate value in response messages.
- o Return Sub-code (1 octet): Set to zero on transmission of request message. The Return Sub-code identifies categories within a specific Return Code and MUST be interpreted within a Return Code.
- o Reserved2 (12 bits): Set to zero on transmission and ignored on reception.
- o F (1 bit): Final flag. When set, indicates this is the last response.
- o C (1 bit): Cross-Connect Error flag (VLAN/FGL mapping error). If set, indicates that the label (VLAN/FGL) in the Flow Entropy is different than the label included in the Diagnostic Label TLV. This field is ignored in request messages and MUST only be interpreted in response messages.
- o O (1 bit): If set, indicates OAM out-of-band response requested.
- o I (1 bit): If set, indicates OAM in-band response requested.

NOTE: When both O and I bits are set to zero, this indicates that no response is required (silent mode). Users MAY specify both O and I, one of them, or none. When both O and I bits are set, the response is sent both in-band and out-of-band.

8.4.4. Out-of-Band Reply Address TLV

The Out-of-Band Reply Address TLV specifies the address to which an out-of-band OAM reply message MUST be sent. When the 0 bit in the TRILL Version sub-TLV (Section 3.3) is not set, the Out-of-Band Reply Address TLV is ignored.




Figure 12: Out-of-Band Reply Address TLV

- o Type (1 octet): 65, Out-of-Band Reply Address TLV
 - o Length (2 octets): Variable. Minimum length is 2 + the length (in octets) of the shortest address. Currently, the minimum value of this field is 4, but this could change in the future if a new address shorter than the TRILL nickname is defined.
 - o Address Type (1 octet):
 - 0 - IPv4
 - 1 - IPv6
 - 2 - TRILL nickname
- All other values reserved.
- o Addr Length (1 octet): Depends on the Address Type. Currently, defined values are:
 - 4 - IPv4
 - 16 - IPv6
 - 2 - TRILL nickname

Other lengths may be acceptable for future Address Types.

- o Reply Address (variable): Address where the reply needs to be sent. Length depends on the address specification.

8.4.5. Diagnostic Label TLV

The Diagnostic Label TLV specifies the data label (VLAN or FGL) in which the OAM messages are generated. Receiving RBridge MUST compare the data label of the Flow Entropy to the data label specified in the Diagnostic Label TLV. The "C" flag (Cross Connect Error) in the response (TRILL OAM Application Identifier TLV; Section 8.4.3) MUST be set when the two VLANs do not match.




Figure 13: Diagnostic Label TLV

- o Type (1 octet): 66, Diagnostic Label TLV
- o Length (2 octets): 5
- o L-Type (1 octet): Label type
 - 0 - Indicates a right-justified 802.1Q 12-bit VLAN padded on the left with bits that must be sent as zero and ignored on receipt
 - 1 - Indicates a TRILL 24-bit fine-grained label
- o Reserved (1 octet): Set to zero on transmission and ignored on reception.
- o Label (24 bits): Either 12-bit VLAN or 24 bit fine-grained label.

RBridges do not perform label error checking when the Diagnostic Label TLV is not included in the OAM message. In certain deployments, intermediate devices may perform label translation. In such scenarios, the originator should not include the Diagnostic Label TLV in OAM messages. Inclusion of Diagnostic Label TLV will generate unwanted label error notifications.

8.4.6. Original Data Payload TLV




Figure 14: Original Data Payload TLV

- o Type (1 octet): 67, Original Data Payload TLV
- o Length (2 octets): variable
- o Original Payload: The original TRILL Header and Flow 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

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

Each TLV can contain up to 255 nicknames of in-scope RBridges. A Multi-destination Tree 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

- o Type (1 octet): 68, RBridge Scope TLV
- o Length (2 octets): Variable. Minimum value is 1.
- o nOfnicknames (1 octet): Indicates the number of nicknames included in this TLV. Zero (0) indicates no nicknames are included in the TLV. When this field is set to zero (0), the Length field MUST be set to 1.
- o Nickname (2 octets): 16-bit RBridge nickname

8.4.8. Previous RBridge Nickname TLV

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

The 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

- o Type (1 octet): 69, Previous RBridge Nickname TLV
- o Length (2 octets): 5

- o Reserved (3 octet): Set to zero on transmission and ignored on reception.
- o Nickname (2 octets): RBridge nickname

8.4.9. Next-Hop RBridge List TLV

The 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.

The 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

- o Type (1 octet): 70, Next-Hop RBridge List TLV
- o Length (2 octets): Variable. Minimum value is 1.
- o nOfnicknames (1 octet): Indicates the number of nicknames included in this TLV. Zero (0) indicates no nicknames are included in the TLV. When this field is set to zero (0), the Length field MUST be set to 1.
- o Nickname (2 octets): 16-bit RBridge nickname

8.4.10. Multicast Receiver Port Count TLV

The 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.

The Multicast Receiver Port Count TLV is an optional TLV.




Figure 18: Multicast Receiver Port Count TLV

- o Type (1 octet): 71, Multicast Receiver Port Count TLV
- o Length (2 octets): 5
- o Reserved (1 octet): Set to zero on transmission and ignored on reception.
- o 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 TLV

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




Figure 19: Flow Identifier TLV

- o Type (1 octet): 72, Flow Identifier TLV
- o Length (2 octets): 5
- o Reserved (1 octet): Set to 0 on transmission and ignored on reception.
- o MEP-ID (2 octets): MEP-ID of the originator [8021Q]. In TRILL, MEP-ID can take a value from 1 to 65535.

- o flow-identifier (2 octets): Uniquely identifies the flow per MEP. Different MEPs may allocate the same flow-identifier value. The {MEP-ID, flow-identifier} pair is globally unique.

Inclusion of the MEP-ID in the Flow Identifier 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

The 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

- o Type (1 octet): 73, Reflector Entropy TLV
- o Length (2 octets): 97
- o Reserved (1 octet): Set to zero on transmission and ignored by the recipient.
- o Reflector Entropy (96 octets): Flow Entropy to be used by the responder. May be padded with zeros if the desired flow-entropy information 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.

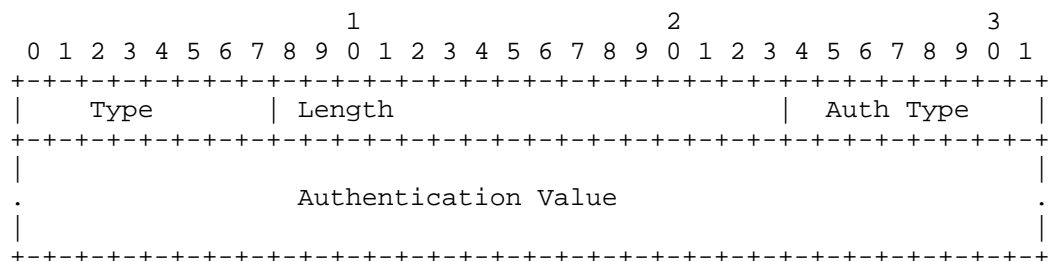


Figure 21: Authentication TLV

- o Type (1 octet): 74, Authentication TLV
- o Length (2 octets): Variable
- o 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. Auth Types 0, 1, 2, and 54 MUST NOT be used. With Auth Type 3, the Authentication TLV is as follows:




Figure 22: Authentication TLV with Auth Type 3

With Auth Type 3, 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, Reserved bit, and Hop Count are treated as zero for the purposes of computing and verifying the Authentication Data.

Key distribution is out of the scope of this document as the keying distributed for IS-IS is used.

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 an Authentication TLV does not verify.

9. Loopback Message

9.1. Loopback Message Format

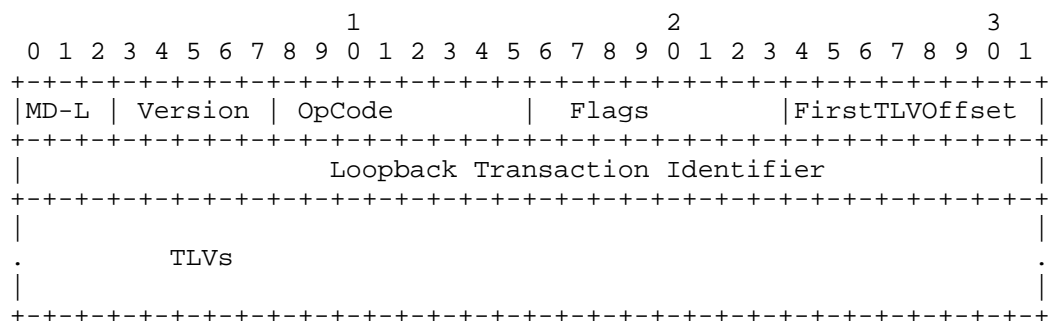


Figure 23: Loopback Message Format

The figure above depicts the format of the Loopback Request and Response messages as defined in [8021Q]. The OpCode for the Loopback Message is set to 3, and the OpCode for the reply message is set to 2 [8021Q]. The Loopback Transaction Identifier (commonly called the Session Identification Number or Session ID in this document) 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:

- o Identifies the destination RBridge nickname based on user specification or based on the specified destination MAC or IP address.
- o Constructs the Flow Entropy based on user-specified parameters or implementation-specific default parameters.
- o Constructs the TRILL OAM header: sets the OpCode to Loopback Message type (3) [8021Q]. Assigns applicable Loopback Transaction Identifier number for the request.
- o The TRILL OAM Application Identifier TLV MUST be included with the flags set to applicable values.
- o Includes following OAM TLVs, where applicable:
 - Out-of-Band Reply Address TLV
 - Diagnostic Label TLV
 - Sender ID TLV
- o Specifies the Hop Count of the TRILL Data frame per user specification or utilize an applicable Hop Count value.
- o Dispatches the OAM frame for transmission.

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

9.2.2. Intermediate RBridge

Intermediate RBridges forward the frame as a normal data frame; 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:

- o The TRILL OAM application encodes the received TRILL Header and Flow Entropy in the Original Data Payload TLV and includes it in the OAM message.
- o 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).
- o Optionally, if the VLAN/FGL identifier value of the received Flow Entropy differs from the value specified in the Diagnostic Label TLV, set the "C" flag (Cross Connect Error) in the TRILL OAM Application Identifier TLV.
- o Include the Sender ID TLV (1).
- o If in-band response was requested, dispatch the frame to the TRILL data plane with request-originator RBridge nickname as the egress RBridge nickname.
- o 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 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 the Loopback Message. The OpCode for Path Trace Reply is 64, and the OpCode for the Path Trace Message is 65.

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 the 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, a Path Trace process timeout occurs, or the Hop Count reaches a configured maximum value.

10.1. Theory of Operation

10.1.1. Actions by Originator RBridge

The originator RBridge takes the following actions:

- o Identifies the destination RBridge based on user specification or based on location of the specified MAC address.
- o Constructs the Flow Entropy based on user-specified parameters or implementation-specific default parameters.
- o Constructs the TRILL OAM header: set the OpCode to Path Trace Message type (65). Assign an applicable Session Identification number for the request. Return Code and Return Sub-code MUST be set to zero.
- o The TRILL OAM Application Identifier TLV MUST be included with the flags set to applicable values.
- o Includes the following OAM TLVs, where applicable:
 - Out-of-Band Reply Address TLV
 - Diagnostic Label TLV
 - Sender ID TLV
- o Specifies the Hop Count of the TRILL Data frame as 1 for the first request.

- o Dispatches 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 retransmission count expires. At each new retransmission, the Session Identification number MUST be incremented. Additionally, for responses received from intermediate R Bridges, the RBridge nickname and interface information MUST be recorded.

10.1.2. Intermediate RBridge

Path Trace Messages transit through Intermediate R Bridges transparently, unless the Hop Count has expired.

The TRILL OAM application layer further validates the received OAM frame by examining the presence of the 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:

- o The TRILL OAM application encodes the received TRILL Header and Flow Entropy in the Original Data Payload TLV and includes it in the OAM message.
- o Set the Return Code to (1) "Reply" and Return Sub-code to two (2) "Intermediate RBridge". Update the TRILL OAM OpCode to 64 (Path Trace Reply).
- o If the VLAN/FGL identifier value of the received Flow Entropy differs from the value specified in the diagnostic label, set the "C" flag (Cross Connect Error) in the TRILL OAM Application Identifier TLV.
- o Include the following TLVs:
 - Previous RBridge Nickname TLV (69)
 - Reply Ingress TLV (5)
 - Reply Egress TLV (6)
 - Interface Status TLV (4)
 - Next-Hop RBridge List TLV (70) (Repeat for each ECMP)
 - Sender ID TLV (1)

- o If a cross-connect error is detected, set the "C" flag (Cross-Connect Error) in the reply's TRILL OAM Application Identifier TLV.
- o If in-band response was requested, dispatch the frame to the TRILL data plane with request-originator RBridge nickname as the egress RBridge nickname.
- o 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 that in Section 10.1.2 with the exception that the TRILL OAM OpCode is set to Path Trace Reply (64).

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 field 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 requests only to RBridges that have failed to respond.

11.1. MTVM Format

The format of MTVM is identical to the Loopback Message format defined in Section 9 with the exception that the OpCode used is 67.

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, the Multicast MAC address and VLAN/FGL, or the 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, and 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 in 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:

- o Identifies the nickname that represents the multicast tree.
- o Obtains the applicable Hop Count value for the selected multicast tree.
- o Constructs TRILL OAM message header and includes the Session Identification number. The Session Identification Number facilitates the originator mapping the response to the correct request.
- o Includes the TRILL OAM Application Identifier TLV, which MUST be included.
- o Includes the OpCode Multicast Tree Verification Message (67).
- o Includes RBridge Scope TLV (68).

- o Optionally, includes the following TLVs, where applicable:
 - Out-of-Band IP Address TLV (65)
 - Diagnostic Label TLV (66)
 - Sender ID TLV (1)
- o Specifies the Hop Count of the TRILL Data frame per user specification or alternatively utilizes the applicable Hop Count value if the TRILL Hop Count is not being specified by the user.
- o Dispatches 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 retransmission count expires. At each new retransmission, the Session Identification Number MUST be incremented. At each retransmission, 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 Section 3.2.

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 the 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:

- o The TRILL OAM application encodes the received TRILL Header and Flow Entropy in the Original Data Payload TLV and includes them in the OAM message.
- o Set the Return Code to zero (0) and Return Sub-code to zero (0). Update the TRILL OAM OpCode to 66 (Multi-destination Tree Verification Reply).

- o Include following TLVs:
 - Previous RBridge Nickname TLV (69)
 - Reply Ingress TLV (5)
 - Interface Status TLV (4)
 - Next-Hop RBridge List TLV (70)
 - Sender ID TLV (1)
 - Multicast Receiver Port Count TLV (71)
- o If a VLAN/FGL cross-connect error is detected, set the "C" flag (Cross-Connect Error) in the TRILL OAM Application Identifier TLV.
- o If in-band response was requested, dispatch the frame to the TRILL data plane with request-originator RBridge nickname as the egress RBridge nickname.
- o 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 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 a 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 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 Identifier TLV, and SNMP Notifications (Traps). Section 12.1 discusses the procedure.

12.1. CCM Error Notification

Each MEP transmits four CCM messages per each flow. ([8021Q] detects CCM fault when three consecutive CCM messages are lost). Each CCM message has a unique sequence number (Session ID) and unique flow-identifier. The flow-identifier is included in the OAM message via the Flow Identifier TLV.

When a 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 an SNMP notification that contains the remote MEP Identification, flow-identifier, and the sequence number of the last CCM message it received, and, if available, the flow-identifier and the sequence number of the first CCM message it received after the failure. Each MEP maintains a unique flow-identifier per each flow; hence, the operator can easily identify flows that correspond to the specific flow-identifier.

The following example illustrates the above.

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

Assume there are three 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-identifier=(1)

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

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

Let's assume Flow-2 is at fault.

MEP-B receives CCM from MEP-A with sequence numbers 1, 2, 3, and 4 but did not receive 5, 6, 7, and 8. CCM timeout is set to three CCM intervals in [8021Q]. Hence, MEP-B detects the error at the 8th CCM message. At this time, the sequence number of the last good CCM message MEP-B has received from MEP-A is 4, and the flow-identifier of the last good CCM Message is (1). Hence, MEP-B will generate a CCM error SNMP notification with MEP-A, last good flow-identifier (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 a CCM message with sequence numbers 9, 10, 11, 12, and 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, the first valid flow-identifier, and the sequence number after the CCM timeout. In the foregoing example, it is MEP-A, flow-identifier (3), and sequence number 9.

The remote MEP list under the CCM MIB Object is augmented to contain "Last Sequence Number", flow-identifier, and "CCM Timeout" variables. "Last Sequence Number" and flow-identifier are updated every time a CCM is received from a remote MEP. The 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

The originator RBridge takes the following actions:

- o Derives the Flow Entropy field based on flow-entropy information specified in the CCM Management object.
- o Constructs the TRILL CCM OAM header as specified in [8021Q].
- o The TRILL OAM Application Identifier TLV MUST be included as the first TLV with the flags set to applicable values.
- o Includes other TLVs specified in [8021Q].
- o Includes the following optional TLV, where applicable:
 - Sender ID TLV (1)
- o Specifies the Hop Count of the TRILL Data frame per user specification or utilize an applicable Hop Count value.

- o Dispatches the OAM frame to the TRILL data plane for transmission.

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

At the 5th retransmission interval, the Flow Entropy of the CCM packet is updated to the next flow-entropy information specified in the CCM Management object. If the 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; no special handling is required.

12.2.3. Destination RBridge

If the CCM Message is addressed to the local RBridge or multicast and satisfies the OAM identification methods specified in Section 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.

The TRILL OAM application layer then validates the MD-Level and pass the packet to the OpCode demultiplexer. The OpCode demultiplexer delivers CCM packets to the CCM process.

The CCM process performs the 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 parts 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 the 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, the process MUST order the fragments based on the Fragment-ID. Any fragments received after the 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 the number of fragments exceeds the maximum supported fragments (255), then the Return Code of the reply message MUST be set to 1 (Reply message), and the Return Sub-code MUST be set to 1 (Fragment limit exceeded).

14. Security Considerations

Forged OAM packets could cause false error or failure indications, mask actual errors or failures, or be used for denial of service. Source addresses for messages can be forged and the out-of-band reply facility (see Section 8.4.4) provides for explicitly supplying the address for replies. For protection against forged OAM packets, the Authentication TLV (see Section 8.4.13) can be used in an OAM message in TRILL. This TLV is virtually identical to the IS-IS Authentication TLV specified in [IS-IS] and depends on IS-IS keying material and the current state of IS-IS keying as discussed in [KARPISIS] and [RFC5310]. In particular, there is currently no standardized IS-IS automated key management.

Of course, authentication is ineffective unless verified and ineffective against senders who have the keying material needed to produce OAM messages that will pass authentication checks. Implementations MUST implement rate-limiting functionality to protect against exploitation of OAM messages as a means of denial-of-service attacks. Aggressive rate-limiting may trigger false positive errors against CCM and LBM-based session monitoring.

Even with authentication, replay of authenticated messages may be possible. There are four types of messages: Continuity Check (CCM), Loopback, Path Trace, and Multi-destination Tree Verification (MTVM). In the case of CCM messages, sequence numbers are required (see Section 12.1) that can protect against replay. In the case of Loopback Messages (see Section 9.1), a Loopback Transaction Identifier is included that, as required by [8021Q], is incremented with each transmission and can detect replays. PTMs (see Section 10) and MTVMs (see Section 11.1) are specified to have the same format as Loopback Messages (although with different OpCodes), so they also have an identifier incremented with each transmission that can detect replays. Thus, all TRILL OAM messages have a field that can be used for replay protection.

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

[8021Q] requires that the MEP filters or passes 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 a 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 framework [RFC4379] 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 has made the assignments described below.

15.1. OAM Capability Flags

Two TRILL-VER sub-TLV Capability Flags (see Section 3.3) have been assigned as follows:

Bit	Description	Reference
---	-----	-----
2	OAM capable	RFC 7455
3	Backwards-compatible OAM	RFC 7455

15.2. CFM Code Points

Four OpCodes have been assigned from the "CFM OAM IETF OpCodes" sub-registry as follows:

Value	Assignment	Reference
-----	-----	-----
64	Path Trace Reply	RFC 7455
65	Path Trace Message	RFC 7455
66	Multi-destination Tree Verification Reply	RFC 7455
67	Multi-destination Tree Verification Message	RFC 7455

Eleven TLV Types have been assigned from the "CFM OAM IETF TLV Types" sub-registry as follows:

Value	Assignment	Reference
-----	-----	-----
64	TRILL OAM Application Identifier TLV	RFC 7455
65	Out-of-Band Reply Address TLV	RFC 7455
66	Diagnostic Label TLV	RFC 7455
67	Original Data Payload TLV	RFC 7455
68	RBridge Scope TLV	RFC 7455
69	Previous RBridge Nickname TLV	RFC 7455
70	Next-Hop RBridge List TLV	RFC 7455
71	Multicast Receiver Port Count TLV	RFC 7455
72	Flow Identifier TLV	RFC 7455
73	Reflector Entropy TLV	RFC 7455
74	Authentication TLV	RFC 7455

15.3. MAC Addresses

IANA has assigned a unicast and a multicast MAC address under the IANA Organizationally Unique Identifier (OUI) for identification of OAM packets as discussed for the backwards-compatibility method (Appendix A.2) and based on the request template in Appendix C. The assigned addresses are 00-00-5E-90-01-00 (unicast) and 01-00-5E-90-01-00 (multicast).

15.4. Return Codes and Sub-codes

IANA has created the "TRILL OAM Return Codes" registry within the "Transparent Interconnection of Lots of Links (TRILL) Parameters" registry and a separate sub-code sub-registry for each Return Code as shown below:

Registry: TRILL OAM Return Codes

Registration Procedure: Standards Action

Return Code	Assignment	References
-----	-----	-----
0	Request message	RFC 7455
1	Reply message	RFC 7455
2-255	Unassigned	RFC 7455

Sub-Registry: Sub-codes for TRILL OAM Return Code 0

Registration Procedure: Standards Action

Sub-code	Assignment	References
-----	-----	-----
0	Valid request	RFC 7455
1-255	Unassigned	RFC 7455

Sub-Registry: Sub-codes for TRILL OAM Return Code 1

Registration Procedure: Standards Action

Sub-code	Assignment	References
-----	-----	-----
0	Valid response	RFC 7455
1	Fragment limit exceeded	RFC 7455
2	Intermediate RBridge	RFC 7455
3-255	Unassigned	RFC 7455

15.5. TRILL Nickname Address Family

IANA has allocated 16396 as the Address Family Number for TRILL nickname.

16. References

16.1. Normative References

- [8021Q] IEEE, "IEEE Standard for Local and metropolitan area networks -- Bridges and Bridged Networks", IEEE Std 802.1Q, December 2014.
- [IS-IS] ISO/IEC, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 2002.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009, <<http://www.rfc-editor.org/info/rfc5310>>.
- [RFC6325] Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011, <<http://www.rfc-editor.org/info/rfc6325>>.
- [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, <<http://www.rfc-editor.org/info/rfc7172>>.

16.2. Informative References

- [KARPISIS] Chunduri, U., Tian, A., and W. Lu, "KARP IS-IS security analysis", Work in Progress, draft-ietf-karp-isis-analysis-04, March 2015.
- [RFC4379] Eronen, P., Ed., and H. Tschofenig, Ed., "Pre-Shared Key Ciphersuites for Transport Layer Security (TLS)", RFC 4279, December 2005, <<http://www.rfc-editor.org/info/rfc4279>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, June 2011, <<http://www.rfc-editor.org/info/rfc6291>>.
- [RFC6361] Carlson, J. and D. Eastlake 3rd, "PPP Transparent Interconnection of Lots of Links (TRILL) Protocol Control Protocol", RFC 6361, August 2011, <<http://www.rfc-editor.org/info/rfc6361>>.
- [RFC6905] Senevirathne, T., Bond, D., Aldrin, S., Li, Y., and R. Watve, "Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)", RFC 6905, March 2013, <<http://www.rfc-editor.org/info/rfc6905>>.
- [RFC7174] Salam, S., Senevirathne, T., Aldrin, S., and D. Eastlake 3rd, "Transparent Interconnection of Lots of Links (TRILL) Operations, Administration, and Maintenance (OAM) Framework", RFC 7174, May 2014, <<http://www.rfc-editor.org/info/rfc7174>>.
- [RFC7176] Eastlake 3rd, 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, <<http://www.rfc-editor.org/info/rfc7176>>.
- [RFC7178] Eastlake 3rd, D., Manral, V., Li, Y., Aldrin, S., and D. Ward, "Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Support", RFC 7178, May 2014, <<http://www.rfc-editor.org/info/rfc7178>>.
- [RFC7179] Eastlake 3rd, D., Ghanwani, A., Manral, V., Li, Y., and C. Bestler, "Transparent Interconnection of Lots of Links (TRILL): Header Extension", RFC 7179, May 2014, <<http://www.rfc-editor.org/info/rfc7179>>.

- [RFC7180] Eastlake 3rd, D., Zhang, M., Ghanwani, A., Manral, V., and A. Banerjee, "Transparent Interconnection of Lots of Links (TRILL): Clarifications, Corrections, and Updates", RFC 7180, May 2014, <<http://www.rfc-editor.org/info/rfc7180>>.
- [RFC7456] Mizrahi, T., Senevirathne, T., Salam, S., Kumar, D., and D. Eastlake 3rd, "Loss and Delay Measurement in Transparent Interconnection of Lots of Links (TRILL)", RFC 7456, March 2015, <<http://www.rfc-editor.org/info/rfc7456>>.
- [TRILLOAMMIB] Kumar, D., Salam, S., and T. Senevirathne, "TRILL OAM MIB", Work in Progress, draft-deepak-trill-oam-mib-01, October 2013.
- [Y1731] ITU-T, "OAM functions and mechanisms for Ethernet based networks", ITU-T Recommendation G.8013/Y.1731, November 2013.

Appendix A. Backwards Compatibility

The methodology presented in this document is in-line with the framework defined in [8021Q] 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 appendix, 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) the Maintenance Point (MEP/MIP) Model and 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 Access Control Lists (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 the TRILL OAM Alert flag presented earlier in this document to identify OAM frames in the hardware.

It is assumed that ECMP path selection of non-IP flows utilizes MAC DA, MAC SA, and VLAN; IP flows utilize IP DA, IP SA, 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 the TRILL Alert flag, as soon as possible, and methods presented here SHOULD be used only as an interim solution.

Identification methods are divided into four 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 Figure 24, 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.

A receiving RBridge identifies OAM flows based on the presence of the well-known values in the specified fields. Additionally, for unicast flows, the egress RBridge nickname of the packet MUST match that of the local RBridge, or for multicast flows, the TRILL Header multicast ("M") flag MUST be set.

Unicast OAM flows that qualify for local processing MUST be redirected to the OAM process and MUST NOT be forwarded (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 the packets MUST be forwarded along the normal path. Additionally, methods MUST be in place to prevent multicast packets from leaking out of the TRILL campus.

Figure 24 summarizes the identification of different OAM frames from data frames.

Flow Entropy	Inner.MacSA	OAM Ethertype	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 24: Identification of TRILL OAM Frames

The unicast and multicast Inner.MacSAs used for the unicast and multicast IP cases, respectively, are 00-00-5E-90-01-00 and 01-00-5E-90-01-00. These have been assigned per the request in Appendix C.

It is important to note that all RBridges MUST generate OAM flows with the "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 MEP-IDs. 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 the 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 Base Mode operation, the value of the Maintenance Domain Name must be the character string "TrillBaseMode" (excluding the quotes). In the Base Mode operation, the Short MA Name format is set to a 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 MEP-ID associated with this MEP is the 2-octet RBridge nickname.

By default, all RBridges operating in Base Mode for TRILL OAM are able to initiate LBM, PTM, 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 25 depicts encoding of MAID within CCM messages.

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

Figure 25: MAID Structure as Defined in [8021Q]

Maintenance Domain Name Format: set to value 4

Maintenance Domain Name Length: set to value 13

Maintenance Domain Name: set to TrillBaseMode

Short MA Name Format: set to value 3

Short MA Name Length: set to value 2

Short MA Name: 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: RFC 7455

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

Acknowledgments

Work on 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, Jari Arkko, Adrian Farrel, Pete Resnick, Stephen Farrell, Dan Romascanu, Gayle Nobel, and Tal Mizrahi.

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

Authors' Addresses

Tissa Senevirathne
Cisco Systems
375 East Tasman Drive
San Jose, CA 95134
United States

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

Norman Finn
Cisco Systems
510 McCarthy Blvd
Milpitas, CA 95035
United States

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
United States

Phone: +1 408-853-9760
EMail: dekumar@cisco.com

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

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

Sam Aldrin
Huawei Technologies
2330 Central Express Way
Santa Clara, CA 95951
United States

EMail: aldrin.ietf@gmail.com

Yizhou Li
Huawei Technologies
101 Software Avenue
Nanjing 210012
China

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

