

TRILL Working Group
Internet Draft
Intended status: Informational

Tissa Senevirathne
CISCO
David Bond
IBM
Sam Aldrin
Yizhou Li
Huawei
Rohit Watve
CISCO

October 20, 2012

Expires: April 2013

**Requirements for Operations, Administration and Maintenance (OAM) in
TRILL
draft-ietf-trill-oam-req-02.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 April 20, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Abstract

OAM (Operations, Administration and Maintenance) is a general term used to identify functions and toolsets to troubleshoot and monitor networks. This document presents, OAM Requirements applicable to TRILL.

Table of Contents

1.	Introduction.....	3
1.1.	Scope.....	3
2.	Conventions used in this document.....	3
3.	Terminology.....	3
4.	OAM Requirements.....	4
4.1.	Data Plane.....	4
4.2.	Connectivity Verification.....	5
4.2.1.	Unicast.....	5
4.2.2.	Multicast.....	5
4.3.	Continuity Check.....	5
4.4.	Path Tracing.....	6
4.5.	General Requirements.....	6
4.6.	Performance Monitoring.....	7
4.6.1.	Packet Loss.....	7
4.6.2.	Packet Delay.....	8
4.7.	ECMP Utilization.....	8
4.8.	Security and Operational considerations.....	8
4.9.	Fault Indications.....	9
4.10.	Defect Indications.....	9
4.11.	Live Traffic monitoring.....	9
5.	Security Considerations.....	10
6.	IANA Considerations.....	10
7.	References.....	10
7.1.	Normative References.....	10
7.2.	Informative References.....	10

8. Acknowledgments.....	11
9. Contributing Authors.....	11

[1. Introduction](#)

OAM (Operations, Administration and Maintenance) generally covers various production aspects of a network. In this document we use the term OAM as defined in [\[RFC6291\]](#).

Success of any mission critical network depends on the ability to proactively monitor networks for faults, performance, etc. as well as its ability to efficiently and quickly troubleshoot defects and failures. A well-defined OAM toolset is a vital requirement for wider adoption of TRILL as the next generation data forwarding technology in larger networks such as data centers.

In this document we define the Requirements for TRILL OAM. It is assumed that the readers are familiar with the OAM concepts and terminologies defined in other OAM standards such as [\[8021ag\]](#), [\[RFC5860\]](#). This document does not attempt to redefine the terms and concepts specified elsewhere.

[1.1. Scope](#)

The scope of this document is OAM between RBridges of a TRILL campus over links selected by TRILL routing.

[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](#) [\[RFC2119\]](#). Although this document is not a protocol specification, the use of this language clarifies the instructions to protocol designers producing solutions that satisfy the requirements set out in this document.

[3. Terminology](#)

Section: The term Section refers to a partial segment of a path between any two given RBridges. As an example, consider the case where RB1 is connected to RBx via RB2, RB3 and RB4. The segment between RB2 to RB4 is referred to as a Section of the path RB1 to RBx.

Flow: The term Flow indicates a set of packets that share the same path and per-hop behavior (such as priority). A flow is typically

identified by a portion of the inner payload that affects the hop-by-hop forwarding decisions. This may contain Layer 2 through Layer 4 information.

All Selectable Least Cost Paths: The term "all selectable least cost paths" refers to a subset of all potentially available least cost paths to a specified destination RBridge that are available (and usable) for forwarding of frames. It is important to note, in practice, due to limitations in implementations, not all available least cost paths may be selectable for forwarding.

Connectivity: The term connectivity indicates reachability between an arbitrary RBridge RB1 and any other RBridge RB2. The specific path can be either explicit (i.e. associated with a specific flow) or unspecified. Unspecified means that messages used for connectivity verification take whatever that path the RBs happen to select.

Continuity Verification: Continuity Verification refers to proactive verification of Connectivity between two RBridges at periodic intervals and generation of explicit notification when Connectivity failures occur.

Fault: The term Fault refers to an inability to perform a required action, e.g., an unsuccessful attempt to deliver a packet.

Defect: The term Defect refers to an interruption in the normal operation, such that over a period of time no packets are delivered successfully.

Failure: The term Failure refers to the termination of the required function over a longer period of time. Persistence of a defect for a period of time is interpreted as a failure.

4. OAM Requirements

4.1. Data Plane

OAM frames, utilized for connectivity verification, continuity checks, performance measurements, etc., will by default take whatever the path TRILL chooses based on the current topology and per-hop equal cost path choices. In some cases, it may be required that the OAM frames utilize specific paths. Thus, it **MUST** be possible to arrange that OAM frames follow the path taken by a specific flow.

RBridges MUST have the ability to identify OAM frames destined for them or which require processing by the OAM plane from normal data frames.

TRILL OAM frames MUST NOT be forwarded out as native frames on end station service enabled ports.

OAM MUST have ability to include all Ethernet traffic types carried by TRILL, including both IP and non-IP traffic.

4.2. Connectivity Verification

4.2.1. Unicast

From an arbitrary RBridge RB1, OAM MUST have the ability to verify connectivity to any other RBridge RB2.

From an arbitrary RBridge RB1, OAM MUST have the ability to verify connectivity to any other RBridge RB2 for a specific flow via the path associated with the specified flow

An RBridge SHOULD have the ability to verify the above connectivity tests on sections. As an example, assume RB1 is connected to RB5 via RB2, RB3 and RB4. An operator SHOULD be able to verify the RB1 to RB5 connectivity on the section from RB3 to RB5. The difference is that the ingress and egress TRILL nicknames in this case are RB1 and RB5 as opposed to RB3 and RB5, even though the message itself may originate at RB3.

4.2.2. Multicast

OAM MUST have the ability to verify connectivity, from an arbitrary RBridge RB1, to either to specific set of RBridges or all member RBridges, for a specified multicast tree. This functionality is referred to as verification of the un-pruned multicast tree.

OAM MUST have the ability to verify connectivity, from an arbitrary RBridge RB1, to either to a specific set of RBridges or all member RBridges, for a specified multicast tree and for a specified flow. This functionality is referred to as verification of the pruned tree.

4.3. Continuity Check

OAM MUST provide functions that allow any arbitrary RBridge RB1 to perform a Continuity Check to any other RBridge.

OAM MUST provide functions that allow any arbitrary RBridge RB1 to perform a Continuity Check to any other RBridge using a path associated with a specified flow.

OAM SHOULD provide functions that allow any arbitrary RBridge to perform a Continuity Check to any other RBridge over all selectable least cost paths.

OAM SHOULD provide the ability to perform a Continuity Check on sections of any selectable path within the network.

OAM SHOULD provide the ability to perform a multicast Continuity Check for specified multi-destination tree(s) as well as specified multi-destination tree and flow combinations. The former is referred to as an un-pruned multi-destination tree Continuity Check and the latter is referred to as a pruned tree Continuity Check.

4.4. Path Tracing

OAM MUST provide the ability to trace a path between any two RBridges per specified unicast flow.

OAM SHOULD provide the ability to trace all selectable least cost paths between any two RBridges.

OAM SHOULD provide functionality to trace all branches of a specified multi-destination tree (un-pruned tree)

OAM SHOULD provide functionality to trace all branches of a specified multi-destination tree for a specified flow (pruned tree).

4.5. General Requirements

OAM MUST provide the ability to initiate and maintain multiple concurrent sessions for multiple OAM functions between any arbitrary RBridge RB1 to any other RBridge. In general, multiple OAM operations will run concurrently. For example, proactive continuity checks may take place between RB1 and RB2 at the same time an operator decides to test connectivity between the same two RBs. Multiple OAM functions and instances of those functions MUST be able to run concurrently without interfering with each other.

OAM MUST provide a single OAM framework for all TRILL OAM functions within the scope of this document.

OAM, as practical and as possible, SHOULD provide a single framework between TRILL and other similar standards.

OAM MUST maintain related error and operational counters. Such counters MUST be accessible via network management applications (e.g. SNMP).

OAM functions related to continuity and connectivity checks MUST be able to be invoked either proactively or on-demand.

OAM SHOULD NOT require extensions to the TRILL header. OAM MAY be required to provide the ability to specify a desired response mode for a specific OAM message. The desired response mode can be either in-band, out-of band or none.

The OAM Framework MUST be extensible to future needs of TRILL and the needs of other standard organizations.

OAM MAY provide methods to verify control plane and forwarding plane alignments.

OAM SHOULD leverage existing OAM technologies, where practical.

4.6. Performance Monitoring

4.6.1. Packet Loss

In this document, term loss of a packet is used as defined in [\[RFC2680\]](#) (see [Section 2.4 of RFC2680](#)).

NOTE: The term simulated flow below indicates a flow that is generated by an RBridge RB1 for OAM purposes. The fields of the simulated flow may or may not be identical to the actual data. However, simulated flow is required to follow the intended path.

OAM SHOULD provide the ability to measure packet loss statistics for a simulated flow from any arbitrary RBridge RB1 to any other RBridge.

OAM SHOULD provide the ability to measure packet loss statistics over a segment, for a simulated flow between any arbitrary RBridge RB1 to any other RBridge.

OAM SHOULD provide the ability to measure simulated packet loss statistics between any two RBridges over all least cost paths.

An RBridge SHOULD be able to perform the above packet loss measurement functions either proactively or on-demand.

4.6.2. Packet Delay

There are two types of packet delays -- one-way delay and two-way delay (Round Trip Delay).

One-way delay is defined in [[RFC2679](#)] as the time elapsed from the start of transmission of the first bit of a packet by an RBridge until the reception of the last bit of the packet by the destination RBridge.

Two-way delay is also referred to as Round Trip Delay is defined similar to [[RFC2681](#)]; i.e. the time elapsed from the start of transmission of the first bit of a packet by an RBridge until the reception of the last bit of the packet by the same RBridge.

OAM SHOULD provide functions to measure two-way delay between two RBridges for a specified flow.

OAM SHOULD provide functions to measure two-way delay between two RBridges for a specified flow over a specific section.

OAM MAY provide functions to measure one-way delay between two RBridges for a specified flow.

OAM MAY provide functions to measure one-way delay between two RBridges for a specified flow over a specific section.

4.7. ECMP Utilization

OAM MAY provide functionality to monitor the effectiveness of per-hop ECMP hashing. For example, individual RBridges could maintain counters that show how packets are being distributed across equal cost next hops for a specified destination RBridge or RBridges as a result of ECMP hashing.

4.8. Security and Operational considerations

Methods MUST be provided to protect against exploitation of OAM framework for security and denial of service attacks.

Methods SHOULD be provided to prevent OAM messages causing congestion in the networks. Periodically generated messages with high frequencies may lead to congestion, hence methods such as shaping or rate limiting SHOULD be utilized.

4.9. Fault Indications

The term Fault refers to an inability to perform a required action, e.g., an unsuccessful attempt to deliver a packet [[OAMOVER](#)]. The unsuccessful attempt may be due to Hop Count expiry, invalid nickname, etc.

OAM MUST provide a Fault Indication framework to notify faults to the ingress RBRidge of the packet or other interested parties (such as syslog servers).

OAM MUST provide functions to selectively enable or disable different types of Fault Indications.

4.10. Defect Indications

[OAMOVER] defines "The term Defect refers to an interruption in the normal operation, such as a consecutive period of time where no packets are delivered successfully."

OAM SHOULD provide a framework for Defect Detection and Indication.

OAM implementations that provide Defect Indication MUST provide methods to selectively enable or disable Defect Detection per defect type.

OAM implementations that provide Defect Indication MUST provide methods to configure Defect Detection thresholds per different types of defects.

OAM implementations that provide Defect Indication facilities MUST provide methods to log defect indications to a locally defined archive such as log buffer or SNMP traps.

OAM implementations that provide Defect Indication facilities SHOULD provide a Remote Defect Indication framework that facilitates notifying the originator/owner of the flow experiencing the defect, which is the ingress RBRidge.

Remote Defect Indication MAY be either in-band or out-of-band.

4.11. Live Traffic monitoring

OAM implementations MAY provide methods to utilize live traffic for troubleshooting and performance monitoring.

Implementations MAY leverage Data Driven CFM [8021Q] or IPFIX [RFC5101] for the purpose of performance monitoring.

5. Security Considerations

Security Requirements are specified in [section 4.8](#). For general TRILL security considerations please refer to [RFC6325]

6. IANA Considerations

None

7. References

7.1. Normative References

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

7.2. Informative References

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

[RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", [RFC5101](#), January 2008.

[RFC2680] Almes, G., et.al. "A One-way Packet Loss Metric for IPPM", [RFC 2680](#), September 1999.

[RFC2679] Almes, G., et.al. "A One-way Delay Metric for IPPM", [RFC 2679](#), September 1999.

[RFC2681] Almes, G., et.al. "A Round-trip Delay Metric for IPPM", [RFC 2681](#), September 1999.

[RFC6291] Anderson, L., et.al. "Guidelines for the Use of the "OAM" Acronym in the IETF", [RFC 6291](#), June 2011.

[8021ag] IEEE, "Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management", 802.1ag, 2007.

[8021Q] IEEE, "Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August, 2011.

- [RFC4377] Nadeau, T., et.al. "Operations and Management (OAM) Requirements for Multi-protocol Label Switched (MPLS) Networks", [RFC 4377](#), February 2006.
- [OAMOVER] Mizrahi, T, et.al., "An Overview of Operations, Administration, and Maintenance (OAM) Mechanisms", [draft-ietf-opsawg-oam-overview-06](#), Work in Progress, March 2012.
- [RFC5860] Vigoureux, M., et.al., "Requirements for Operations, Administration and Maintenance (OAM) in MPLS Transport Networks", [RFC5860](#), May 2010.

8. Acknowledgments

Special acknowledgments to IEEE 802.1 chair, Tony Jeffree for allowing us to solicit comments from IEEE 802.1 group. Also recognized are the comments received from IEEE group, Ayal Lior and others.

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

9. Contributing Authors

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

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

David Bond
IBM
2051 Mission College Blvd
Santa Clara, CA 95054
USA

Phone: +1-603-339-7575
Email: mokon@mokon.net

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

Rohit Watve
CISCO Systems
375 East Tasman Drive
San Jose, CA 95134
USA.

Phone: +1-408-424-2091
Email: rwatve@cisco.com

Thomas Narten
IBM Corporation
3039 Cornwallis Avenue,
PO Box 12195
Research Triangle Park, NC 27709
USA

Email: narten@us.ibm.com

Donald Eastlake
Huawei Technologies
155 Beaver Street,
Milford, MA 01757
USA.

Email: d3e3e3@gmail.com

Anoop Ghanwani
DELL
350 Holger Way
San Jose, CA 95134
USA.

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

Jon Hudson
Brocade
120 Holger Way
San Jose, CA 95134
USA.

Email: jon.hudson@gmail.com

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

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

Radia Perlman
Intel Labs
2700 156th Ave NE, Suite 300,
Bellevue, WA 98007
USA.

Phone: +1-425-881-4824
Email: radia.perlman@intel.com

Tal Mizrahi
Marvell
6 Hamada St.
Yokneam, 20692 Israel

Email: talmi@marvell.com