

Internet Engineering Task Force (IETF)
Request for Comments: 9062
Category: Informational
ISSN: 2070-1721

S. Salam
A. Sajassi
Cisco
S. Aldrin
Google
J. Drake
Juniper
D. Eastlake 3rd
Futurewei
June 2021

Framework and Requirements for Ethernet VPN (EVPN) Operations, Administration, and Maintenance (OAM)

Abstract

This document specifies the requirements and reference framework for Ethernet VPN (EVPN) Operations, Administration, and Maintenance (OAM). The requirements cover the OAM aspects of EVPN and Provider Backbone Bridge EVPN (PBB-EVPN). The framework defines the layered OAM model encompassing the EVPN service layer, network layer, underlying Packet Switched Network (PSN) transport layer, and link layer but focuses on the service and network layers.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9062>.

Copyright Notice

Copyright (c) 2021 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
 - 1.1. Relationship to Other OAM Work
 - 1.2. Specification of Requirements
 - 1.3. Terminology

- 2. EVPN OAM Framework
 - 2.1. OAM Layering
 - 2.2. EVPN Service OAM
 - 2.3. EVPN Network OAM
 - 2.4. Transport OAM for EVPN
 - 2.5. Link OAM
 - 2.6. OAM Interworking
- 3. EVPN OAM Requirements
 - 3.1. Fault Management Requirements
 - 3.1.1. Proactive Fault Management Functions
 - 3.1.1.1. Fault Detection (Continuity Check)
 - 3.1.1.2. Defect Indication
 - 3.1.1.2.1. Forward Defect Indication (FDI)
 - 3.1.1.2.2. Reverse Defect Indication (RDI)
 - 3.1.2. On-Demand Fault Management Functions
 - 3.1.2.1. Connectivity Verification
 - 3.1.2.2. Fault Isolation
 - 3.2. Performance Management
 - 3.2.1. Packet Loss
 - 3.2.2. Packet Delay and Jitter
- 4. Security Considerations
- 5. IANA Considerations
- 6. References
 - 6.1. Normative References
 - 6.2. Informative References

Acknowledgements

Authors' Addresses

1. Introduction

This document specifies the requirements and defines a reference framework for Ethernet VPN (EVPN) Operations, Administration, and Maintenance (OAM) [RFC6291]. In this context, we use the term "EVPN OAM" to loosely refer to the OAM functions required for and/or applicable to [RFC7432] and [RFC7623].

EVPN is a Layer 2 VPN (L2VPN) solution for multipoint Ethernet services with advanced multihoming capabilities that uses BGP for distributing Customer/Client Media Access Control (C-MAC) address reachability information over the core MPLS/IP network.

PBB-EVPN combines Provider Backbone Bridging (PBB) [IEEE-802.1Q] with EVPN in order to reduce the number of BGP MAC advertisement routes; provide client MAC address mobility using C-MAC [RFC7623] aggregation and Backbone MAC (B-MAC) [RFC7623] sub-netting; confine the scope of C-MAC learning to only active flows; offer per-site policies; and avoid C-MAC address flushing on topology changes.

This document focuses on the fault management and performance management aspects of EVPN OAM. It defines the layered OAM model encompassing the EVPN service layer, network layer, underlying Packet Switched Network (PSN) transport layer, and link layer but focuses on the service and network layers.

1.1. Relationship to Other OAM Work

This document leverages concepts and draws upon elements defined and/or used in the following documents:

[RFC6136] specifies the requirements and a reference model for OAM as it relates to L2VPN services, pseudowires, and associated Packet Switched Network (PSN) tunnels. This document focuses on Virtual Private LAN Service (VPLS) and Virtual Private Wire Service (VPWS) solutions and services.

[RFC8029] defines mechanisms for detecting data plane failures in

MPLS Label Switched Paths (LSPs), including procedures to check the correct operation of the data plane as well as mechanisms to verify the data plane against the control plane.

[IEEE-802.1Q] specifies the Ethernet Connectivity Fault Management (CFM) protocol, which defines the concepts of Maintenance Domains, Maintenance Associations, Maintenance End Points, and Maintenance Intermediate Points.

[Y.1731] extends Connectivity Fault Management in the following areas: it defines fault notification and alarm suppression functions for Ethernet and specifies mechanisms for Ethernet performance management, including loss, delay, jitter, and throughput measurement.

1.2. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.3. Terminology

This document uses the following terminology, much of which is defined in [RFC6136]:

CE	Customer Edge device; for example, a host, router, or switch.
CFM	Connectivity Fault Management [IEEE-802.1Q]
DF	Designated Forwarder [RFC7432]
Down MEP	A MEP that originates traffic away from and terminates traffic towards the core of the device in whose port it is logically located.
EVI	An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN [RFC7432].
L2VPN	Layer 2 VPN
LOC	Loss of continuity
MA	Maintenance Association; a set of MEPs belonging to the same Maintenance Domain (MD) established to verify the integrity of a single service instance [IEEE-802.1Q].
MD	Maintenance Domain; an OAM Domain that represents a region over which OAM frames can operate unobstructed [IEEE-802.1Q].
MEP	Maintenance End Point; it is responsible for origination and termination of OAM frames for a given MA. A MEP is logically located in a device's port [IEEE-802.1Q].
MIP	Maintenance Intermediate Point; it is located between peer MEPs and can process and respond to certain OAM frames but does not initiate them. A MIP is logically located in a device's port [IEEE-802.1Q].
MP2P	Multipoint to Point
NMS	Network Management Station [RFC6632]

P	Provider network interior (non-edge) node
P2MP	Point to Multipoint
PBB	Provider Backbone Bridge [RFC7623]
PE	Provider Edge network device
Up MEP	A MEP that originates traffic towards and terminates traffic from the core of the device in whose port it is logically located.
VPN	Virtual Private Network

2. EVPN OAM Framework

2.1. OAM Layering

Multiple layers come into play for implementing an L2VPN service using the EVPN family of solutions as listed below. The focus of this document is the service and network layers.

- * The service layer runs end to end between the sites or Ethernet segments that are being interconnected by the EVPN solution.
- * The network layer extends between the EVPN PE (Provider Edge) nodes and is mostly transparent to the P (provider network interior) nodes (except where flow entropy comes into play). It leverages MPLS for service (i.e., EVI) multiplexing and split-horizon functions.
- * The transport layer is dictated by the networking technology of the PSN. It may be based on either MPLS LSPs or IP.
- * The link layer is dependent upon the physical technology used. Ethernet is a popular choice for this layer, but other alternatives are deployed (e.g., Packet over SONET (POS), Dense Wavelength Division Multiplexing (DWDM), etc.).

This layering extends to the set of OAM protocols that are involved in the ongoing maintenance and diagnostics of EVPN networks. Figure 1 below depicts the OAM layering and shows which devices have visibility into what OAM layer(s).

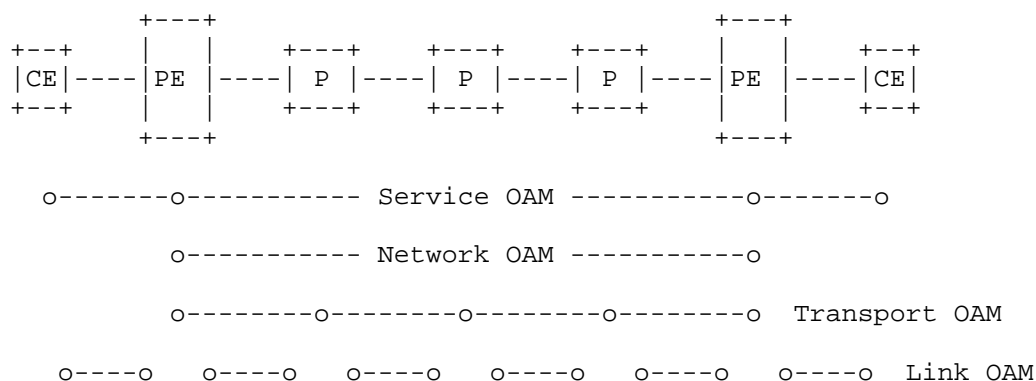


Figure 1: OAM Layering

Service OAM and Network OAM mechanisms only have visibility to the PE nodes but not the P nodes. As such, they can be used to deduce whether the fault is in the customer's own network, the local CE-PE segment, the PE-PE segment, or the remote CE-PE segment(s). EVPN Transport OAM mechanisms can be used for fault isolation between the

PEs and P nodes.

Figure 2 below shows an example network where Ethernet domains are interconnected via EVPN using MPLS, and it shows the OAM mechanisms that are applicable at each layer. The details of the layers are described in the sections below.

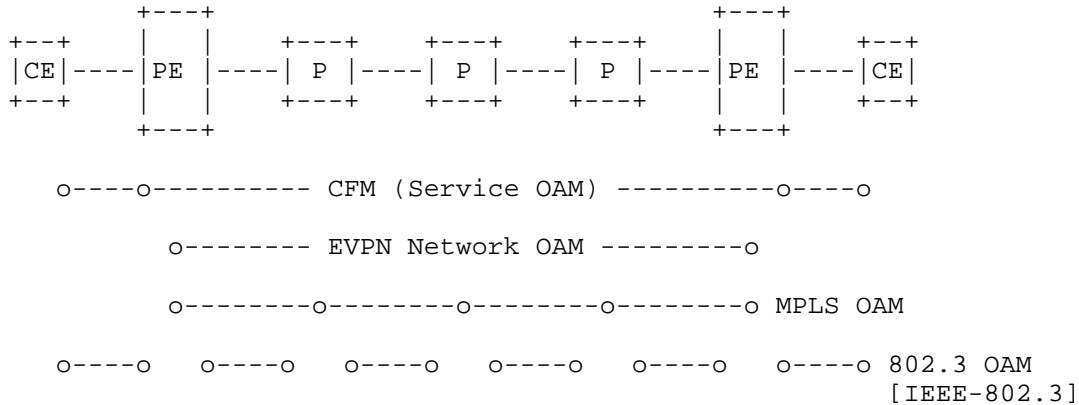


Figure 2: EVPN OAM Example

2.2. EVPN Service OAM

The EVPN Service OAM protocol depends on what service-layer technology is being interconnected by the EVPN solution. In the case of [RFC7432] and [RFC7623], the service layer is Ethernet; hence, the corresponding Service OAM protocol is Ethernet CFM [IEEE-802.1Q].

EVPN Service OAM is visible to the CEs and EVPN PEs but not to the P nodes. This is because the PEs operate at the Ethernet MAC layer in [RFC7432] and [RFC7623], whereas the P nodes do not.

The EVPN PE MUST support MIP functions in the applicable Service OAM protocol (for example, Ethernet CFM). The EVPN PE SHOULD support MEP functions in the applicable Service OAM protocol. This includes both Up and Down MEP functions.

As shown in Figure 3, the MIP and MEP functions being referred to are logically located within the device's port operating at the customer level. (There could be MEPs/MIPs within PE ports facing the provider network, but they would not be relevant to EVPN Service OAM as the traffic passing through them will be encapsulated/tunneled, so any customer-level OAM messages will just be treated as data.) Down MEP functions are away from the core of the device (towards the PE forwarding mechanism in the case of a PE). OAM messages between the PE Up MEPs shown are a type of EVPN Network OAM, while such messages between the CEs or from a PE to its local CE or to the remote CE are Service OAMs.




Figure 3: CFM Details

The EVPN PE MUST, by default, learn the MAC address of locally attached CE MEPs by snooping on CFM frames and advertising them to remote PEs as a MAC/IP Advertisement route. Some means to limit the number of MAC addresses that a PE will learn SHOULD be implemented.

The EVPN PE SHOULD advertise any MEP/MIP local to the PE as a MAC/IP Advertisement route. Since these are not subject to mobility, they SHOULD be advertised with the static (sticky) bit set (see Section 15.2 of [RFC7432]).

2.3. EVPN Network OAM

EVPN Network OAM is visible to the PE nodes only. This OAM layer is analogous to Virtual Circuit Connectivity Verification (VCCV) [RFC5085] in the case of VPLS/VPWS. It provides mechanisms to check the correct operation of the data plane as well as a mechanism to verify the data plane against the control plane. This includes the ability to perform fault detection and diagnostics on:

- * the MP2P tunnels used for the transport of unicast traffic between PEs. EVPN allows for three different models of unicast label assignment: label per EVI, label per <ESI, Ethernet Tag>, and label per MAC address. In all three models, the label is bound to an EVPN Unicast Forwarding Equivalence Class (FEC). EVPN Network OAM MUST provide mechanisms to check the operation of the data plane and verify that operation against the control plane view.
- * the MP2P tunnels used for aliasing unicast traffic destined to a multihomed Ethernet segment. The three label assignment models, discussed above, apply here as well. In all three models, the label is bound to an EVPN Aliasing FEC. EVPN Network OAM MUST provide mechanisms to check the operation of the data plane and verify that operation against the control plane view.
- * the multicast tunnels (either MP2P or P2MP) used for the transport of broadcast, unknown unicast, and multicast traffic between PEs. In the case of ingress replication, a label is allocated per EVI or per <EVI, Ethernet Tag> and is bound to an EVPN Multicast FEC. In the case of Label Switched Multicast (LSM) and, more specifically, aggregate inclusive trees, again, a label may be allocated per EVI or per <EVI, Ethernet Tag> and is bound to the tunnel FEC.
- * the correct operation of the Ethernet Segment Identifier (ESI) split-horizon filtering function. In EVPN, a label is allocated per multihomed Ethernet segment for the purpose of performing the access split-horizon enforcement. The label is bound to an EVPN Ethernet segment.
- * the correct operation of the Designated Forwarder (DF) [RFC7432] filtering function. EVPN Network OAM MUST provide mechanisms to check the operation of the data plane and verify that operation against the control plane view for the DF filtering function.

EVPN Network OAM mechanisms MUST provide in-band monitoring capabilities. It is desirable, to the extent practical, for OAM test messages to share fate with data messages. Details of how to achieve this are beyond the scope of this document.

EVPN Network OAM SHOULD provide both proactive and on-demand mechanisms of monitoring the data plane operation and data plane conformance to the state of the control plane.

2.4. Transport OAM for EVPN

The Transport OAM protocol depends on the nature of the underlying transport technology in the PSN. MPLS OAM mechanisms [RFC8029] [RFC6425] as well as ICMP [RFC0792] and ICMPv6 [RFC4443] are applicable, depending on whether the PSN employs MPLS or IP transport, respectively. Furthermore, Bidirectional Forwarding Detection (BFD) mechanisms per [RFC5880], [RFC5881], [RFC5883], and [RFC5884] apply. Also, the BFD mechanisms pertaining to MPLS-TP LSPs per [RFC6428] are applicable.

2.5. Link OAM

Link OAM depends on the data-link technology being used between the PE and P nodes. For example, if Ethernet links are employed, then Ethernet Link OAM ([IEEE-802.3], Clause 57) may be used.

2.6. OAM Interworking

When interworking two networking domains, such as actual Ethernet and EVPN to provide an end-to-end emulated service, there is a need to identify the failure domain and location, even when a PE supports both the Service OAM mechanisms and the EVPN Network OAM mechanisms. In addition, scalability constraints may not allow the running of proactive monitoring, such as Ethernet Continuity Check Messages (CCMs) [IEEE-802.1Q], at a PE to detect the failure of an EVI across the EVPN domain. Thus, the mapping of alarms generated upon failure detection in one domain (e.g., actual Ethernet or EVPN network domain) to the other domain is needed. There are also cases where a PE may not be able to process Service OAM messages received from a remote PE over the PSN even when such messages are defined, as in the Ethernet case, thereby necessitating support for fault notification message mapping between the EVPN Network domain and the Service domain.

OAM interworking is not limited, though, to scenarios involving disparate network domains. It is possible to perform OAM interworking across different layers in the same network domain. In general, alarms generated within an OAM layer, as a result of proactive fault detection mechanisms, may be injected into its client-layer OAM mechanisms. This allows the client-layer OAM to trigger event-driven (i.e., asynchronous) fault notifications. For example, alarms generated by the Link OAM mechanisms may be injected into the Transport OAM layer, and alarms generated by the Transport OAM mechanism may be injected into the Network OAM mechanism, and so on.

EVPN OAM MUST support interworking between the Network OAM and Service OAM mechanisms. EVPN OAM MAY support interworking among other OAM layers.

3. EVPN OAM Requirements

This section discusses the EVPN OAM requirements pertaining to fault management and performance management.

3.1. Fault Management Requirements

3.1.1. Proactive Fault Management Functions

The network operator configures proactive fault management functions to run periodically. Certain actions (for example, protection switchover or alarm indication signaling) can be associated with specific events, such as entering or clearing fault states.

3.1.1.1. Fault Detection (Continuity Check)

Proactive fault detection is performed by periodically monitoring the reachability between service end points, i.e., MEPs in a given MA, through the exchange of CCMs [IEEE-802.1Q]. The reachability between any two arbitrary MEPs may be monitored for:

- * in-band, per-flow monitoring. This enables per-flow monitoring between MEPs. EVPN Network OAM MUST support fault detection with per-user flow granularity. EVPN Service OAM MAY support fault detection with per-user flow granularity.
- * a representative path. This enables a liveness check of the nodes hosting the MEPs, assuming that the loss of continuity (LOC) to the MEP is interpreted as a failure of the hosting node. This, however, does not conclusively indicate liveness of the path(s) taken by user data traffic. This enables node failure detection but not path failure detection through the use of a test flow. EVPN Network OAM and Service OAM MUST support fault detection using test flows.
- * all paths. For MPLS/IP networks with ECMP, the monitoring of all unicast paths between MEPs (on non-adjacent nodes) may not be possible since the per-hop ECMP hashing behavior may yield situations where it is impossible for a MEP to pick flow entropy characteristics that result in exercising the exhaustive set of ECMP paths. The monitoring of all ECMP paths between MEPs (on non-adjacent nodes) is not a requirement for EVPN OAM.

The fact that MPLS/IP networks do not enforce congruency between unicast and multicast paths means that the proactive fault detection mechanisms for EVPN networks MUST provide procedures to monitor the unicast paths independently of the multicast paths. This applies to EVPN Service OAM and Network OAM.

3.1.1.2. Defect Indication

Defect indications can be categorized into two types: forward and reverse, as described below. EVPN Service OAM MUST support at least one of these types of event-driven defect indications upon the detection of a connectivity defect.

3.1.1.2.1. Forward Defect Indication (FDI)

FDI is used to signal a failure that is detected by a lower-layer OAM mechanism. A server MEP (i.e., an actual or virtual MEP) transmits a forward defect indication in a direction away from the direction of the failure (refer to Figure 4 below).




Figure 4: Forward Defect Indication

Forward defect indication may be used for alarm suppression and/or for the purpose of interworking with other layer OAM protocols. Alarm suppression is useful when a transport-level or network-level fault translates to multiple service- or flow-level faults. In such a scenario, it is enough to alert a network management station (NMS)

of the single transport-level or network-level fault in lieu of flooding that NMS with a multitude of Service or Flow granularity alarms. EVPN PEs SHOULD support forward defect indication in the Service OAM mechanisms.

3.1.1.2.2. Reverse Defect Indication (RDI)

RDI is used to signal that the advertising MEP has detected a LOC defect. RDI is transmitted in the direction of the failure (refer to Figure 5).




Figure 5: Reverse Defect Indication

RDI allows single-sided management, where the network operator can examine the state of a single MEP and deduce the overall health of a monitored service. EVPN PEs SHOULD support reverse defect indication in the Service OAM mechanisms. This includes both the ability to signal a LOC defect to a remote MEP as well as the ability to recognize RDI from a remote MEP. Note that, in a multipoint MA, RDI is not a useful indicator of unidirectional fault. This is because RDI carries no indication of the affected MEP(s) with which the sender had detected a LOC defect.

3.1.2. On-Demand Fault Management Functions

On-demand fault management functions are initiated manually by the network operator and continue for a bounded time period. These functions enable the operator to run diagnostics to investigate a defect condition.

3.1.2.1. Connectivity Verification

EVPN Network OAM MUST support on-demand connectivity verification mechanisms for unicast and multicast destinations. The connectivity verification mechanisms SHOULD provide a means for specifying and carrying the following in the messages:

- * variable-length payload/padding to test connectivity problems related to the Maximum Transmission Unit (MTU).
- * test frame formats as defined in Appendix C of [RFC2544] to detect potential packet corruption.

EVPN Network OAM MUST support connectivity verification at per-flow granularity. This includes both user flows (to test a specific path between PEs) as well as test flows (to test a representative path between PEs).

EVPN Service OAM MUST support connectivity verification on test flows and MAY support connectivity verification on user flows.

For multicast connectivity verification, EVPN Network OAM MUST support reporting on:

- * the DF filtering status of a specific port(s) or all the ports in

a given bridge domain.

- * the split-horizon filtering status of a specific port(s) or all the ports in a given bridge domain.

3.1.2.2. Fault Isolation

EVPN OAM MUST support an on-demand fault localization function. This involves the capability to narrow down the locality of a fault to a particular port, link, or node. The characteristic of forward/reverse path asymmetry in MPLS/IP makes fault isolation a direction-sensitive operation. That is, given two PEs A and B, localization of continuity failures between them requires running fault-isolation procedures from PE A to PE B as well as from PE B to PE A.

EVPN Service OAM mechanisms only have visibility to the PEs but not the MPLS or IP P nodes. As such, they can be used to deduce whether the fault is in the customer's own network, the local CE-PE segment, or a remote CE-PE segment(s). EVPN Network and Transport OAM mechanisms can be used for fault isolation between the PEs and P nodes.

3.2. Performance Management

Performance management functions can be performed both proactively and on demand. Proactive management involves a recurring function, where the performance management probes are run continuously without a trigger. We cover both proactive and on-demand functions in this section.

3.2.1. Packet Loss

EVPN Network OAM SHOULD provide mechanisms for measuring packet loss for a given service -- for example, [RFC7680] and [RFC6673].

Given that EVPN provides inherent support for multipoint-to-multipoint connectivity, packet loss cannot be accurately measured by means of counting user data packets. This is because user packets can be delivered to more PEs or more ports than are necessary (e.g., due to broadcast, unpruned multicast, or unknown unicast flooding). As such, a statistical means of approximating the packet loss rate is required. This can be achieved by sending "synthetic" OAM packets that are counted only by those ports (MEPs) that are required to receive them. This provides a statistical approximation of the number of data frames lost, even with multipoint-to-multipoint connectivity.

3.2.2. Packet Delay and Jitter

EVPN Service OAM SHOULD support measurement of one-way and two-way packet delay and delay variation (jitter) across the EVPN network. Measurement of one-way delay requires clock synchronization between the probe source and target devices. Mechanisms for clock synchronization are outside the scope of this document. Note that Service OAM performance management mechanisms defined in [Y.1731] can be used. See also [RFC7679], [RFC2681], and [RFC3393].

EVPN Network OAM MAY support measurement of one-way and two-way packet delay and delay variation (jitter) across the EVPN network.

4. Security Considerations

EVPN OAM MUST prevent OAM packets from leaking outside of the EVPN network or outside their corresponding Maintenance Domain. This can be done for CFM, for example, by having MEPs implement a filtering function based on the Maintenance Level associated with received OAM

packets.

EVPN OAM SHOULD provide mechanisms for implementation and optional use to:

- * prevent denial-of-service attacks caused by exploitation of the OAM message channel (for example, by forging messages to exceed a Maintenance End Point's capacity to maintain state).
- * authenticate communicating end points (for example, MEPS and MIPS).

5. IANA Considerations

This document has no IANA actions.

6. References

6.1. Normative References

- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <<https://www.rfc-editor.org/info/rfc5881>>.
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <<https://www.rfc-editor.org/info/rfc5883>>.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, DOI 10.17487/RFC5884, June 2010, <<https://www.rfc-editor.org/info/rfc5884>>.
- [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, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/info/rfc6291>>.
- [RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", RFC 6425, DOI 10.17487/RFC6425, November 2011, <<https://www.rfc-editor.org/info/rfc6425>>.
- [RFC6428] Allan, D., Ed., Swallow, G., Ed., and J. Drake, Ed.,

"Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, DOI 10.17487/RFC6428, November 2011, <<https://www.rfc-editor.org/info/rfc6428>>.

- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7623] Sajassi, A., Ed., Salam, S., Bitar, N., Isaac, A., and W. Henderickx, "Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)", RFC 7623, DOI 10.17487/RFC7623, September 2015, <<https://www.rfc-editor.org/info/rfc7623>>.
- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

6.2. Informative References

- [IEEE-802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE Std 802.1Q-2014, DOI 10.1109/IEEESTD.2014.6991462, December 2014, <<https://doi.org/10.1109/IEEESTD.2014.6991462>>.
- [IEEE-802.3] IEEE, "IEEE Standard for Ethernet", IEEE Std 802.3-2018, DOI 10.1109/IEEESTD.2018.8457469, August 2018, <<https://doi.org/10.1109/IEEESTD.2018.8457469>>.
- [RFC2544] Bradner, S. and J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, DOI 10.17487/RFC2544, March 1999, <<https://www.rfc-editor.org/info/rfc2544>>.
- [RFC2681] Almes, G., Kalidindi, S., and M. Zekauskas, "A Round-trip Delay Metric for IPPM", RFC 2681, DOI 10.17487/RFC2681, September 1999, <<https://www.rfc-editor.org/info/rfc2681>>.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", RFC 3393, DOI 10.17487/RFC3393, November 2002, <<https://www.rfc-editor.org/info/rfc3393>>.
- [RFC5085] Nadeau, T., Ed. and C. Pignataro, Ed., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, DOI 10.17487/RFC5085, December 2007, <<https://www.rfc-editor.org/info/rfc5085>>.
- [RFC6136] Sajassi, A., Ed. and D. Mohan, Ed., "Layer 2 Virtual Private Network (L2VPN) Operations, Administration, and Maintenance (OAM) Requirements and Framework", RFC 6136, DOI 10.17487/RFC6136, March 2011, <<https://www.rfc-editor.org/info/rfc6136>>.
- [RFC6632] Ersue, M., Ed. and B. Claise, "An Overview of the IETF Network Management Standards", RFC 6632, DOI 10.17487/RFC6632, June 2012,

<<https://www.rfc-editor.org/info/rfc6632>>.

- [RFC6673] Morton, A., "Round-Trip Packet Loss Metrics", RFC 6673, DOI 10.17487/RFC6673, August 2012, <<https://www.rfc-editor.org/info/rfc6673>>.
- [RFC7679] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Delay Metric for IP Performance Metrics (IPPM)", STD 81, RFC 7679, DOI 10.17487/RFC7679, January 2016, <<https://www.rfc-editor.org/info/rfc7679>>.
- [RFC7680] Almes, G., Kalidindi, S., Zekauskas, M., and A. Morton, Ed., "A One-Way Loss Metric for IP Performance Metrics (IPPM)", STD 82, RFC 7680, DOI 10.17487/RFC7680, January 2016, <<https://www.rfc-editor.org/info/rfc7680>>.
- [Y.1731] ITU-T, "Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks", ITU-T Recommendation G.8013/Y.1731, August 2015.

Acknowledgements

The authors would like to thank the following for their review of this work and their valuable comments: David Black, Martin Duke, Xiao Min, Gregory Mirsky, Zaheduzzaman Sarker, Dave Schinazi, John Scudder, Melinda Shore, Robert Wilton, Alexander Vainshtein, Stig Venaas, and ric Vyncke.

Authors' Addresses

Samer Salam
Cisco
The Atrium Building, Floor 3
Weygand St.
Beirut
Lebanon

Email: ssalam@cisco.com

Ali Sajassi
Cisco
170 West Tasman Drive
San Jose, CA 95134
United States of America

Email: sajassi@cisco.com

Sam Aldrin
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
United States of America

Email: aldrin.ietf@gmail.com

John E. Drake
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States of America

Email: jdrake@juniper.net

Donald E. Eastlake 3rd
Futurewei Technologies
2386 Panoramic Circle
Apopka, FL 32703
United States of America

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