Routing Area Working Group
Internet-Draft
Intended status: Informational
Expires: January 9, 2017

G. Mirsky
Ericsson
E. Nordmark
Arista Networks
C. Pignataro
N. Kumar
D. Kumar
Cisco Systems, Inc.
M. Chen
Y. Li
Huawei Technologies
D. Mozes
Mellanox Technologies Ltd.
S. Pallagatti

I. Bagdonas
July 8, 2016

Operations, Administration and Maintenance (OAM) for Overlay Networks: Gap Analysis
draft-ooamdt-rtgwg-oam-gap-analysis-02

Abstract

This document provides an overview of the Operations, Administration, and Maintenance (OAM) for overlay networks. The OAM toolset includes set of fault management and performance monitoring capabilities (operating in the data plane) that comply with the Overlay OAM Requirements. Insufficient functional coverage of existing OAM protocols also noted in this document. The protocol definitions for each of the Overlay OAM tools to be defined in separate documents.

This Internet-Draft will expire on January 9, 2017.

Table of Contents

1. Introduction

Operations, Administration, and Maintenance (OAM) toolset provides methods for fault management and performance monitoring in each layer of the network, in order to improve their ability to support services with guaranteed and strict Service Level Agreements (SLAs) while reducing operational costs.

[RFC7276] provided detailed analysis of OAM protocols. Since its completion several new protocols that define data plane encapsulation were introduced. That presented both need to re-evaluate existing set of OAM tools and opportunity to build it into set of tools that can be used and re-used for different data plane protocols.

[I-D.ooamdt-rtgwg-ooam-requirement] defines the set of requirements for OAM in Overlay networks. The OAM solution for Overlay networks, developed by the design team, has two objectives:

- The Overlay OAM toolset should be developed based on existing IP and IP/MPLS architecture, technology, and toolsets.

- The Overlay OAM operational experience should be similar to that in other, e.g. IP and IP/MPLS, networks.

1.1. Conventions used in this document

1.1.1. Terminology

Term "Overlay OAM" used in this document interchangeably with longer version "set of OAM protocols, methods and tools for Overlay networks".

AIS Alarm Indication Signal

BFD Bidirectional Forwarding Detection

BIER Bit-Indexed Explicit Replication

CC Continuity Check

CV Connectivity Verification

FM Fault Management

G-ACh Generic Associated Channel

Geneve Generic Network Virtualization Encapsulation

GUE Generic UDP Encapsulation

MPLS Multiprotocol Label Switching

NTP Network Time Protocol

NVO3 Network Virtalization Overlays

OAM Operations, Administration, and Maintenance

OWAMP One-Way Active Measurement Protocol

PM Performance Measurement

PTP Precision Time Protocol

SFC Service Fundction Chaining

SFP Service Function Path

SLA Service Level Agreement

TWAMP Two-Way Active Measurement Protocol

VxLAN Virtual eXtensible Local Area Network

VxLAN-GPE Generic Protocol Extension for VxLAN

1.1.2. Requirements Language

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 [RFC2119].

2. Working Group Overview

2.1. BIER

The BIER working group has some WG documents on OAM which are discussed further in this document.

2.2. NVO3

The NVO3 encapsulations (Geneve [I-D.ietf-nvo3-geneve], GUE [I-D.ietf-nvo3-gue], and GPE [I-D.ietf-nvo3-vxlan-gpe]) all have some notion of a OAM bit or flag. In Geneve this is defined to not apply to intermediate (underlay) routers and that the setting of the bit

doesn't affect the ECMP hash. The other proposals do not have as succinct constraints on their OAM bit/flag.

There are currently no NVO3 working group OAM protocol specifications. The OAM documents that have been discussed are individual drafts such as [I-D.ashwood-nvo3-oam-requirements], [I-D.nordmark-nvo3-transcending-traceroute], [I-D.pang-nvo3-vxlan-path-detection], [I-D.saum-nvo3-pmtud-over-vxlan], and [I-D.singh-nvo3-vxlan-router-alert].

2.3. SFC

TBD

3. Overlay OAM Toolset

It is expected that the encapsulation of an overlay network uses one of methods discussed in [I-D.ietf-rtgwg-dt-encap] to distinctly identify the payload as OAM, i.e. non-user, packet. In its turn all Overlay OAM protocols share the common Overlay OAM Header. Format and processing of the header are outside the scope of this document and will be presented in the solution document.

3.1. Overlay OAM Fault Management

Protocols that enable Fault Management functions of OAM toolset are comprised of protocols that perform proactive and on-demand defect detection and failure localization.

3.1.1. Proactive Continuity Check and Connectivity Verification

Bidirectional Forwarding Detection (BFD) has been designed as proactive Continuity Check protocol. [RFC6428] defined extension to support Connectivity Verification in MPLS-TP networks . Following BFD specifications can be used in overlay networks:

- BFD for point-to-point as defined in [RFC5880], [RFC5882], [RFC5883], [RFC5884], [RFC5885], [RFC6428] and [RFC7726];
- BFD for multipoint network as defined in [I-D.ietf-bfd-multipoint] and [I-D.ietf-bfd-multipoint-active-tail];
- S-BFD as defined in [I-D.ietf-bfd-seamless-base] and [I-D.ietf-bfd-seamless-ip];

3.1.1.1. Proactive CC/CV in BIER

. Bit-Indexed Explicit Replication (BIER) provides the multicast service. For that BFD over multipoint network [I-D.ietf-bfd-multipoint] and [I-D.ietf-bfd-multipoint-active-tail] are the most suitable of BFD family Figure 1 presents IP/UDP format of BFD over BIER in MPLS network.

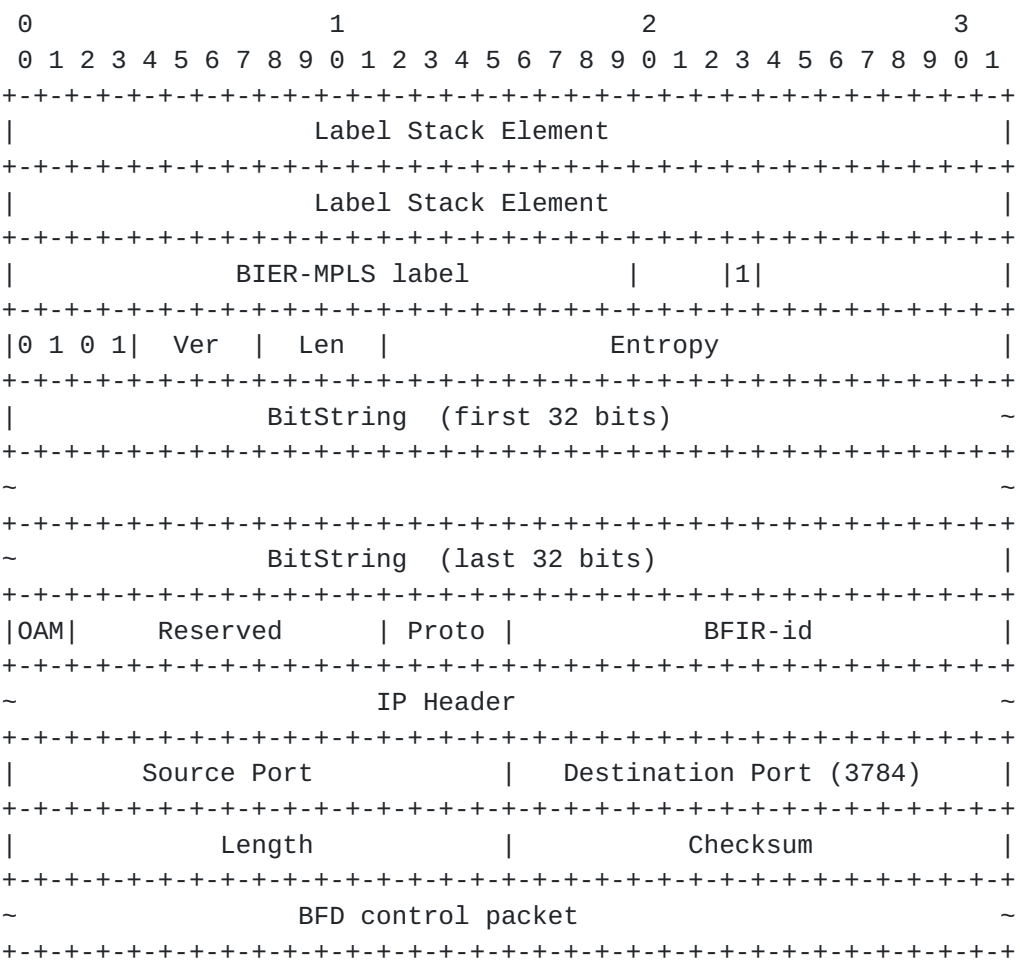

Figure 1: BFD over BIER with IP/UDP format

Proto field MUST be set to IPv4 or IPv6 vlalue. Note that IP Destination address in Figure 1 must follow Section 7 [RFC5884], i.e. ?the destination IP address MUST be randomly chosen from the 127/8 range for IPv4 and from the 0:0:0:0:0:FFFF:7F00/104 range for IPv6.? BFD packets in the reverse direction of the BFD session will be transmitted on IP network to the IP address mapped to the BFIR-id and the destination UDP port number set as source UDP port number of the received BFD packet.

IP/UDP format presents overhead, particularly in case of IPv6 address family. Thus option to avoid use of extra headers for OAM seems attractive. Figure 2 presents G-ACh format of BFD over BIER in MPLS network. Proto field of the BIER header MUST be set to OAM value. BFD control packet follows the BIER OAM header as defined in [I-D.kumarzheng-bier-ping]. According to the Section 3.1 of [I-D.kumarzheng-bier-ping], Ver is set to 1; BFD control packet over multi-point without or with active tail accordingly identified in Message Type Field. The Proto field ?is used to define if there is any data packet immediately following the OAM payload?.

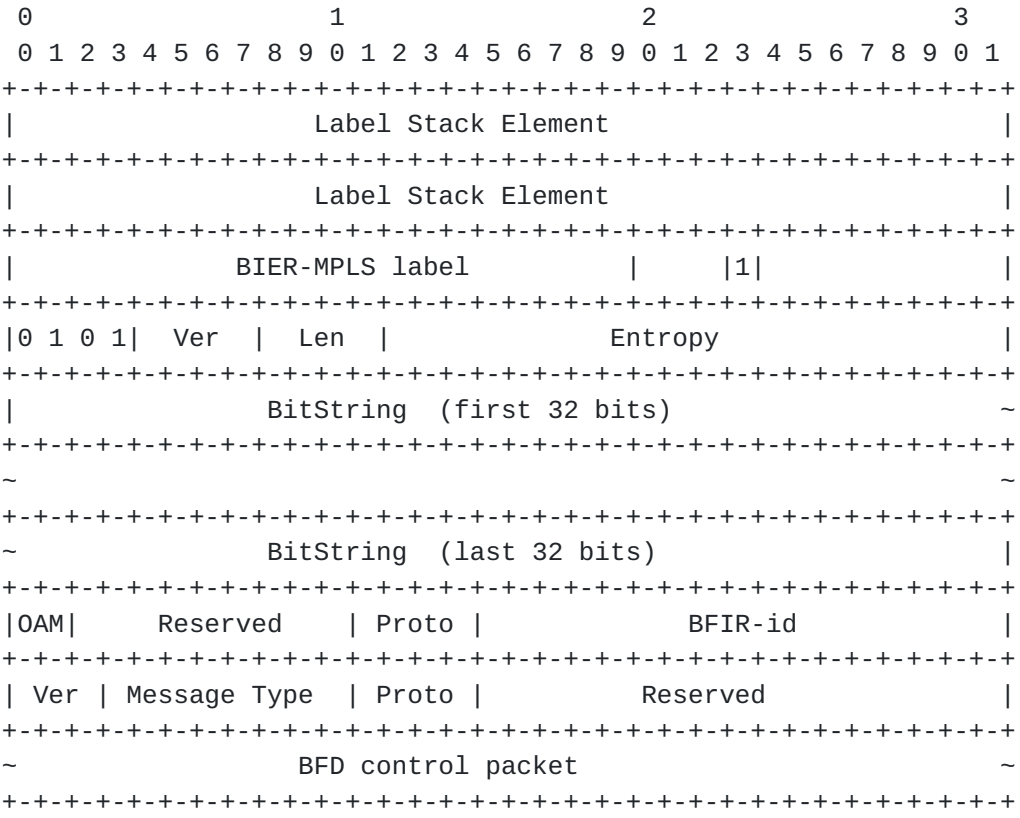

Figure 2: BFD over BIER with G-ACh format

3.1.1.2. Proactive CC/CV in NVO3

There is currently no WG document on proactive CC/CV. The individual requirements document [I-D.ashwood-nvo3-oam-requirements] covers this and there is a related proposal for BFD over VXLAN in [I-D.spallagatti-bfd-vxlan].

3.1.1.3. Proactive CC/CV over SFP

3.1.2. On-demand Continuity Check and Connectivity Verification

On-demand Continuity Check and Connectivity Verification protocols include:

- MPLS Echo Request/Reply, a.k.a. LSP Ping, as defined in [RFC4379] and its numerous extensions;
- LSP Self-ping, as defined in [RFC7746];
- [I-D.kumarzheng-bier-ping] is a good example of generic troubleshooting and defect localization tool that can be extended and suited for more specific requirements of the particular type of an overlay network;

3.1.2.1. On-demand CC/CV in BIER

[I-D.kumarzheng-bier-ping] defines format of Echo Request/Reply control packet and set of TLVs that can be used to perform failure detection and isolation in BIER domain over MPLS network.

3.1.2.2. On-demand CC/CV in NVO3

There is currently no WG document for on-demand CC/CV.

Individual documents exist for tracing such as [I-D.pang-nvo3-vxlan-path-detection], and [I-D.nordmark-nvo3-transcending-traceroute].

3.1.2.3. On-demand CC/CV over SFP

3.1.3. Alarm Indication Signal

3.1.3.1. AIS in BIER

3.1.3.2. AIS in NVO3

There is currently no WG document on Alarm Indication Signal.

The individual draft [I-D.nordmark-nvo3-transcending-traceroute] suggests reusing ICMP errors for defect indications.

3.1.3.3. AIS over SFP

3.2. Overlay OAM Performance Measurement

These protocols may be considered for Overlay Performance Measurement (PM) OAM:

- packet loss and delay measurement in MPLS networks, as defined in [RFC6374] with ability to export measurement results for post-processing [I-D.ietf-mpls-rfc6374-udp-return-path];

- One-Way Active Measurement Protocol (OWAMP), as defined in [RFC4656], and Two-Way Active Measurement Protocol (TWAMP), as defined in [RFC5357], [RFC6038], and [RFC7750];

- use of the Marking Method [I-D.ietf-ippm-alt-mark] that, if accordingly supported by the overlay layer, can behave as close as technically possible to a passive method to measure performance, e.g. [I-D.mirsky-bier-pmmm-oam].

3.2.1. Overlay OAM PM Active

Requirements towards PM OAM for overlay networks are listed in the Section 4.2 [I-D.ooamdt-rtgwg-ooam-requirement]. Two sets of performance measurement protocols had been developed at IETF so far:

- OWAMP [RFC4656] and TWAMP [RFC5357] each includes the control protocol to negotiate required parameters and control a test session as well as the test protocol itself that specify format and processing of a test packet. Historically, TWAMP, that enables measurement of the latency, packet loss both as one-way and round trip performance metric, has been implemented more often and thus gained wider deployment than OWAMP. There are several properties of the test protocol that may not be suitable for its use in overlay networks:

  - the test protocol is targetted to IP layer and carries some IP-specific information;

  - the format of the sent test and the reflected packets differ significantly and that complicates efficient HW-based implementation;

  - latency and packet loss measurement operations are bundled together and that causes certain overhead when only one of performance metrics is to be measured;

- only Network Time Protocol (NTP) format of timestamp is currently supported that requires additional processing to convert from IEEE-1588 time format that broadly supported in many current packet forwarding engines.

o [RFC6374] defines the test protocol that enables measurement of the latency and paket loss as one-way and round-trip perfomance metrics. Comparing to OWAMP/TWAMP RFC 6374 has certain advantages:

- the test protocol is flexible and these performance metrics can be measured independently or in the single test session;

- the protocol does not carry transport layer specific information;

- there's no difference between format of the packet transmitted by the sender and reflected by the responder as the test packets preallocates space for all necesary data it collects;

- both NTP and PTP time formats allowed to be used to record time for latency measurement.

[RFC6374] can be used as foundation of active PM OAM in overlay networks. The YANG data model [RFC6020] of the packet loss and delay measurement based on [RFC6374] can improve control and increase operational value of active performance measurement in overlay networks.

3.2.1.1. Active PM in BIER

Currently there is no draft related to active PM OAM in the WG.

3.2.1.2. Active PM in NVO3

Performance management has been discussed in NVO3 but there is currently no draft in the WG.

3.2.1.3. Active PM over SFP

3.2.2. Overlay OAM PM Passive

3.2.2.1. Passive PM in BIER

[I-D.mirsky-bier-pmmm-oam] describes how the Marking Method can be used in BIER domain over MPLS networks.

3.2.2.2. Passive PM in NVO3

Marking has been discussed in NVO3 sessions, but there is no draft in the working group.

The Generic Protocol Extension for VXLAN [I-D.ietf-nvo3-vxlan-gpe], Generic Network Virtualization Encapsulation [I-D.ietf-nvo3-geneve], Generic UDP Encapsulation [I-D.ietf-nvo3-gue] are just some examples of the new encapsulations to support network virtualization. NVO3 PM would be used to probe the NV Edge to NV Edge tunnels and NV Edge entity status for a DC network. The main requirement for Performance Management is to be able to support measurement of the frame loss, delay and delay variation between two NV Edge devices that support the same VNI within a given NVO3 domain on per VNI basis. Alternate Marking Method [I-D.ietf-ippm-alt-mark] enables calculation of these metrics but sets forth requirements toward overlay encapsulation to make use of the AMM behave in the network as passive OAM per definition in [RFC7799].

3.2.2.3. Passive PM over SFP

In the SFC architecture SF, SFF, Classifier and NSH Proxy Agent are the elements that can incorporate the measurement agent functionality to support SFC performance measurement. The required OAM Performance Measurement, as described in [I-D.ietf-sfc-oam-framework] highlight the capability to assess the monitoring at SF and SFF or a Set of SF/SFF, both in case of SFC-aware SF and SFC-unaware SF; the monitoring of SFP (and RSP) that comprises a set of SFs that may be ordered or unordered; the monitoring of the Classifiers operation and the monitoring of the SFC as a whole.

Performance measurement includes measuring of packet loss, delay, delay variation and could be performed by the marking method proposed in [I-D.ietf-ippm-alt-mark]. To make use of the marking method behave as passive OAM, as defined in [RFC7799], the overlay network encapsulation should allocate the field, preferrably two bits long, whose value does not affect how a packet is treated by the overlay network.

3.3. Telemetry in Overlay OAM

Excessive use of the in-band OAM channel may affect user flow and thus change network behavior. For example, if operator uses passive measurement exporting massive amount of data over the OAM channel may affect network. I think that a management channel should be used in such case. Obviously it may traverse the same nodes and links but may not require the same QoS. We can refer to LMAP Reference Model [RFC7594] with Controller, Measurement Agent and Data Collector.

[I-D.lapukhov-dataplane-probe] proposes transport independent generic telemetry probe structure.

3.4. Conclusions

- A common Overlay OAM header should be defined to support demultiplexing of OAM protocols.
- Existing modes of BFD protocol, primarily its Async mode, can be used either in IP/UDP or ACH format, as proactive continuity check mechanism in overlay networks.
- A new control packet to be used for on-demand CC/CV in overlay networks should be defined. Set of common TLVs may be defined while more specific TLVs to be defined by respective groups of experts.
- [RFC6374] can be used as the foundation of active performance measurement OAM in overlay networks.
- YANG data model of the active PM OAM in overlay networks would improve control and increase operational value of the test methods.
- Performance measurement includes measuring of packet loss, delay, delay variation and could be performed by the marking method, for example as proposed in [I-D.ietf-ippm-alt-mark]. To make use of the marking method behave as passive OAM, as defined in [RFC7799], the overlay network encapsulation should allocate the field, preferrably two bits long, whose value does not affect how a packet is treated by the overlay network.

4. IANA Considerations

This document does not propose any IANA consideration. This section may be removed.

5. Security Considerations

6. Acknowledgement

TBD

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

7.2. Informative References

[I-D.ashwood-nvo3-oam-requirements]
Chen, H., Ashwood-Smith, P., Xia, L., Iyengar, R., Tsou, T., Sajassi, A., Boucadair, M., Jacquenet, C., Daikoku, M., Ghanwani, A., and R. Krishnan, "NVO3 Operations, Administration, and Maintenance Requirements", draft-ashwood-nvo3-oam-requirements-04 (work in progress), October 2015.

[I-D.ietf-bfd-multipoint]
Katz, D., Ward, D., and J. Networks, "BFD for Multipoint Networks", draft-ietf-bfd-multipoint-08 (work in progress), April 2016.

[I-D.ietf-bfd-multipoint-active-tail]
Katz, D., Ward, D., and J. Networks, "BFD Multipoint Active Tails.", draft-ietf-bfd-multipoint-active-tail-02 (work in progress), May 2016.

[I-D.ietf-bfd-seamless-base]
Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and J. Networks, "Seamless Bidirectional Forwarding Detection (S-BFD)", draft-ietf-bfd-seamless-base-11 (work in progress), May 2016.

[I-D.ietf-bfd-seamless-ip]
Pignataro, C., Ward, D., and N. Akiya, "Seamless Bidirectional Forwarding Detection (S-BFD) for IPv4, IPv6 and MPLS", draft-ietf-bfd-seamless-ip-06 (work in progress), May 2016.

[I-D.ietf-ippm-alt-mark]
Fioccola, G., Capello, A., Cociglio, M., Castaldelli, L., Chen, M., Zheng, L., Mirsky, G., and T. Mizrahi, "Alternate Marking method for passive performance monitoring", draft-ietf-ippm-alt-mark-01 (work in progress), July 2016.

[I-D.ietf-mpls-rfc6374-udp-return-path]
Bryant, S., Sivabalan, S., and S. Soni, "RFC6374 UDP Return Path", draft-ietf-mpls-rfc6374-udp-return-path-05 (work in progress), April 2016.

[I-D.ietf-nvo3-geneve]
Gross, J. and I. Ganga, "Geneve: Generic Network Virtualization Encapsulation", draft-ietf-nvo3-geneve-01 (work in progress), January 2016.

[I-D.ietf-nvo3-gue]
Herbert, T., Yong, L., and O. Zia, "Generic UDP Encapsulation", draft-ietf-nvo3-gue-04 (work in progress), July 2016.

[I-D.ietf-nvo3-vxlan-gpe]
Kreeger, L. and U. Elzur, "Generic Protocol Extension for VXLAN", draft-ietf-nvo3-vxlan-gpe-02 (work in progress), April 2016.

[I-D.ietf-rtgwg-dt-encap]
Nordmark, E., Tian, A., Gross, J., Hudson, J., Kreeger, L., Garg, P., Thaler, P., and T. Herbert, "Encapsulation Considerations", draft-ietf-rtgwg-dt-encap-01 (work in progress), March 2016.

[I-D.ietf-sfc-oam-framework]
Aldrin, S., Krishnan, R., Akiya, N., Pignataro, C., and A. Ghanwani, "Service Function Chaining Operation, Administration and Maintenance Framework", draft-ietf-sfc-oam-framework-01 (work in progress), February 2016.

[I-D.kumarzheng-bier-ping]
Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", draft-kumarzheng-bier-ping-03 (work in progress), July 2016.

[I-D.lapukhov-dataplane-probe]
Lapukhov, P. and r. remy@barefootnetworks.com, "Data-plane probe for in-band telemetry collection", draft-lapukhov-dataplane-probe-01 (work in progress), June 2016.

[I-D.mirsky-bier-pmmm-oam]
Mirsky, G., Zheng, L., Chen, M., and G. Fioccola, "Performance Measurement (PM) with Marking Method in Bit Index Explicit Replication (BIER) Layer", draft-mirsky-bier-pmmm-oam-01 (work in progress), March 2016.

[I-D.nordmark-nvo3-transcending-traceroute]
Nordmark, E., Appanna, C., and A. Lo, "Layer-Transcending Traceroute for Overlay Networks like VXLAN", draft-nordmark-nvo3-transcending-traceroute-02 (work in progress), March 2016.

[I-D.ooamdt-rtgwg-ooam-requirement]
Kumar, N., Pignataro, C., Kumar, D., Mirsky, G., Chen, M., Nordmark, E., Networks, J., and D. Mozes, "Overlay OAM Requirements", draft-ooamdt-rtgwg-ooam-requirement-01 (work in progress), July 2016.

[I-D.pang-nvo3-vxlan-path-detection]
<>, J., Liu, D., Liu, D., Zhang, D., Yizhou, L., Chen, H., <>, D., <>, B., Kumar, D., Gao, R., and Y. Qiao, "Path Detection in VXLAN Overlay Network", draft-pang-nvo3-vxlan-path-detection-02 (work in progress), March 2016.

[I-D.saum-nvo3-pmtud-over-vxlan]
Dikshit, S. and A. Nayak, "PMTUD Over Vxlan", draft-saum-nvo3-pmtud-over-vxlan-03 (work in progress), June 2016.

[I-D.singh-nvo3-vxlan-router-alert]
Singh, K., Jain, P., Rio, D., Henderickx, W., Shekhar, R., and R. Rahman, "VxLAN Router Alert Option", draft-singh-nvo3-vxlan-router-alert-02 (work in progress), September 2015.

[I-D.spallagatti-bfd-vxlan]
Networks, J., sajibasil@gmail.com, s., Paragiri, S., Govindan, V., Mudigonda, M., and G. Mirsky, "BFD for VXLAN", draft-spallagatti-bfd-vxlan-03 (work in progress), April 2016.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, DOI 10.17487/RFC4379, February 2006, <http://www.rfc-editor.org/info/rfc4379>.

[RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, DOI 10.17487/RFC4656, September 2006, <http://www.rfc-editor.org/info/rfc4656>.

[RFC5357] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and J. Babiarz, "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, DOI 10.17487/RFC5357, October 2008, <http://www.rfc-editor.org/info/rfc5357>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <http://www.rfc-editor.org/info/rfc5880>.

[RFC5882] Katz, D. and D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, DOI 10.17487/RFC5882, June 2010, <http://www.rfc-editor.org/info/rfc5882>.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <http://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, <http://www.rfc-editor.org/info/rfc5884>.

[RFC5885] Nadeau, T., Ed. and C. Pignataro, Ed., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, DOI 10.17487/RFC5885, June 2010, <http://www.rfc-editor.org/info/rfc5885>.

[RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <http://www.rfc-editor.org/info/rfc6020>.

[RFC6038] Morton, A. and L. Ciavattone, "Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features", RFC 6038, DOI 10.17487/RFC6038, October 2010, <http://www.rfc-editor.org/info/rfc6038>.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <http://www.rfc-editor.org/info/rfc6374>.

[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, <http://www.rfc-editor.org/info/rfc6428>.

[RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <http://www.rfc-editor.org/info/rfc7276>.

[RFC7594] Eardley, P., Morton, A., Bagnulo, M., Burbridge, T., Aitken, P., and A. Akhter, "A Framework for Large-Scale Measurement of Broadband Performance (LMAP)", RFC 7594, DOI 10.17487/RFC7594, September 2015, <http://www.rfc-editor.org/info/rfc7594>.

[RFC7726] Govindan, V., Rajaraman, K., Mirsky, G., Akiya, N., and S. Aldrin, "Clarifying Procedures for Establishing BFD Sessions for MPLS Label Switched Paths (LSPs)", RFC 7726, DOI 10.17487/RFC7726, January 2016, <http://www.rfc-editor.org/info/rfc7726>.

[RFC7746] Bonica, R., Minei, I., Conn, M., Pacella, D., and L. Tomotaki, "Label Switched Path (LSP) Self-Ping", RFC 7746, DOI 10.17487/RFC7746, January 2016, <http://www.rfc-editor.org/info/rfc7746>.

[RFC7750] Hedin, J., Mirsky, G., and S. Baillargeon, "Differentiated Service Code Point and Explicit Congestion Notification Monitoring in the Two-Way Active Measurement Protocol (TWAMP)", RFC 7750, DOI 10.17487/RFC7750, February 2016, <http://www.rfc-editor.org/info/rfc7750>.

[RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <http://www.rfc-editor.org/info/rfc7799>.

Authors' Addresses

Greg Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Erik Nordmark
Arista Networks

Email: nordmark@acm.org

Carlos Pignataro
Cisco Systems, Inc.

Email: cpignata@cisco.com

Nagendra Kumar
Cisco Systems, Inc.

Email: naikumar@cisco.com

Deepak Kumar
Cisco Systems, Inc.

Email: dekumar@cisco.com

Mach Chen
Huawei Technologies

Email: mach.chen@huawei.com

Yizhou Li
Huawei Technologies

Email: liyizhou@huawei.com

David Mozes
Mellanox Technologies Ltd.

Email: davidm@mellanox.com

Santosh Pallagatti

Email: santosh.pallagatti@gmail.com

Ignas Bagdonas

Email: ibagdona@gmail.com