

DetNet Working Group
Internet-Draft
Intended status: Standards Track
Expires: 15 July 2024

G. Mirsky
Ericsson
M. Chen
Huawei
B. Varga
Ericsson
12 January 2024

Operations, Administration and Maintenance (OAM) for Deterministic
Networks (DetNet) with MPLS Data Plane
draft-ietf-detnet-mpls-oam-15

Abstract

This document defines format and usage principles of the Deterministic Network (DetNet) service Associated Channel over a DetNet network with the MPLS data plane. The DetNet service Associated Channel can be used to carry test packets of active Operations, Administration, and Maintenance protocols that are used to detect DetNet failures and measure performance metrics.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on 15 July 2024.

Copyright Notice

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions used in this document	3
2.1. Terminology and Acronyms	3
2.2. Keywords	4
3. Active OAM for DetNet Networks with MPLS Data Plane	4
3.1. DetNet Active OAM Encapsulation	4
3.2. DetNet Packet Replication, Elimination, and Ordering Functions Interaction with Active OAM	7
4. OAM Interworking Models	8
4.1. OAM of DetNet MPLS Interworking with OAM of TSN	8
4.2. OAM of DetNet MPLS Interworking with OAM of DetNet IP	9
5. IANA Considerations	10
5.1. DetNet Associated Channel Header Flags Registry	10
6. Security Considerations	10
7. Acknowledgment	10
8. References	10
8.1. Normative References	10
8.2. Informational References	11
Authors' Addresses	13

1. Introduction

[RFC8655] introduces and explains Deterministic Networks (DetNet) architecture and how the Packet Replication, Elimination, and Ordering functions (PREOF) can be used to ensure a low packet drop ratio in a DetNet domain.

Operations, Administration, and Maintenance (OAM) protocols are used to detect and localize network defects, and to monitor network performance. Some OAM functions (e.g., failure detection) are usually performed proactively in the network, while others (e.g., defect localization) are typically performed on demand. These tasks can be achieved through a combination of active and hybrid OAM methods, as classified in [RFC7799]. This document presents a format for active OAM in DetNet networks with MPLS data plane.

Also, this document defines format and usage principles of the DetNet service Associated Channel over a DetNet network with the MPLS data plane [RFC8964].

2. Conventions used in this document

2.1. Terminology and Acronyms

The term "DetNet OAM" is used in this document interchangeably with longer version "set of OAM protocols, methods and tools for Deterministic Networks".

DetNet Deterministic Network

d-ACH DetNet Associated Channel Header

OAM Operations, Administration, and Maintenance

PREOF Packet Replication, Elimination, and Ordering Functions

PW Pseudowire

E2E End-to-end

BFD Bidirectional Forwarding Detection

TSN IEEE 802.1 Time-Sensitive Networking

CFM Connectivity Fault Management

F-Label - a DetNet "forwarding" label. The F-Label identifies the LSP used to forward a DetNet flow across an MPLS PSN, e.g., a hop-by-hop label used between label switching routers.

S-Label - a DetNet "service" label. An S-Label is used between DetNet nodes that implement the DetNet service sub-layer functions. An S-Label is also used to identify a DetNet flow at DetNet service sub-layer.

Underlay Network or Underlay Layer - the network that provides connectivity between the DetNet nodes. One example of an underlay layer is an MPLS network that provides Label Switched Path (LSP) connectivity between DetNet nodes.

DetNet Node - a node that is an actor in the DetNet domain. Examples of DetNet nodes include DetNet domain Edge nodes, and DetNet nodes that perform PREOF within the DetNet domain.

2.2. Keywords

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.

3. Active OAM for DetNet Networks with MPLS Data Plane

OAM protocols and mechanisms act within the data plane of the particular networking layer, thus it is critical that the data plane encapsulation supports OAM mechanisms that comply with the OAM requirements listed in [I-D.ietf-detnet-oam-framework].

Operation of a DetNet data plane with an MPLS underlay network is specified in [RFC8964]. Within the MPLS underlay network, DetNet flows are to be encapsulated analogous to pseudowires as specified in [RFC3985], [RFC4385]. For reference, the Generic Pseudowire (PW) MPLS Control Word (as defined in [RFC4385] and used with DetNet) is reproduced in Figure 1.

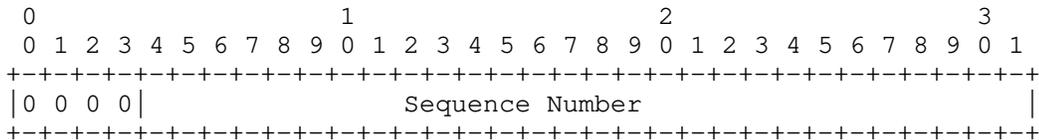


Figure 1: DetNet Control Word Format

PREOF in the DetNet domain is composed of a combination of nodes that perform replication and elimination functions. The Elimination sub-function always uses the S-Label in conjunction with the packet sequencing information (i.e., the Sequence Number encoded in the DetNet Control Word). The Replication sub-function uses the S-Label information only.

3.1. DetNet Active OAM Encapsulation

DetNet OAM, like PW OAM, uses the PW Associated Channel Header defined in [RFC4385]. At the same time, a DetNet PW can be viewed as a Multi-Segment PW, where DetNet service sub-layer functions are at the segment endpoints. However, DetNet service sub-layer functions operate per packet level (not per segment). These per-packet level characteristics of PREOF require additional fields for proper OAM packet processing. Encapsulation of a DetNet MPLS [RFC8964] active

OAM packet is shown in Figure 2.

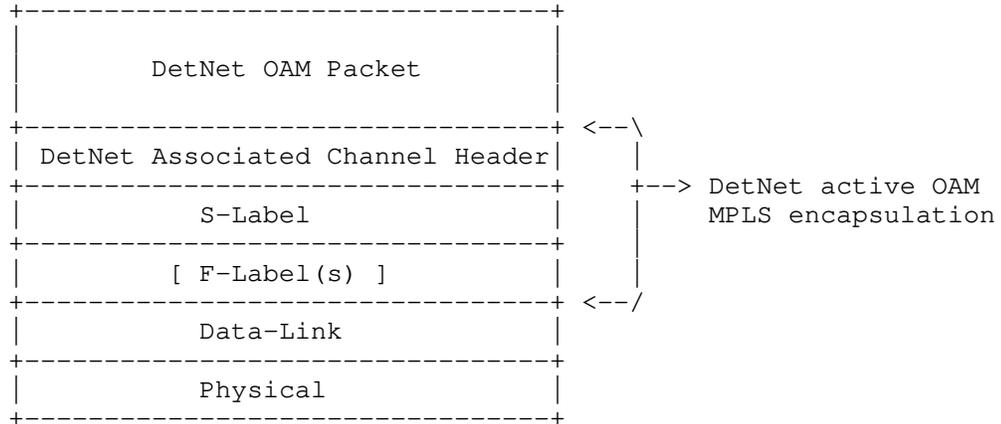


Figure 2: DetNet Active OAM Packet Encapsulation in MPLS Data Plane

Figure 3 displays encapsulation of a test packet of an active DetNet OAM protocol in case of MPLS-over-UDP/IP [RFC9025].

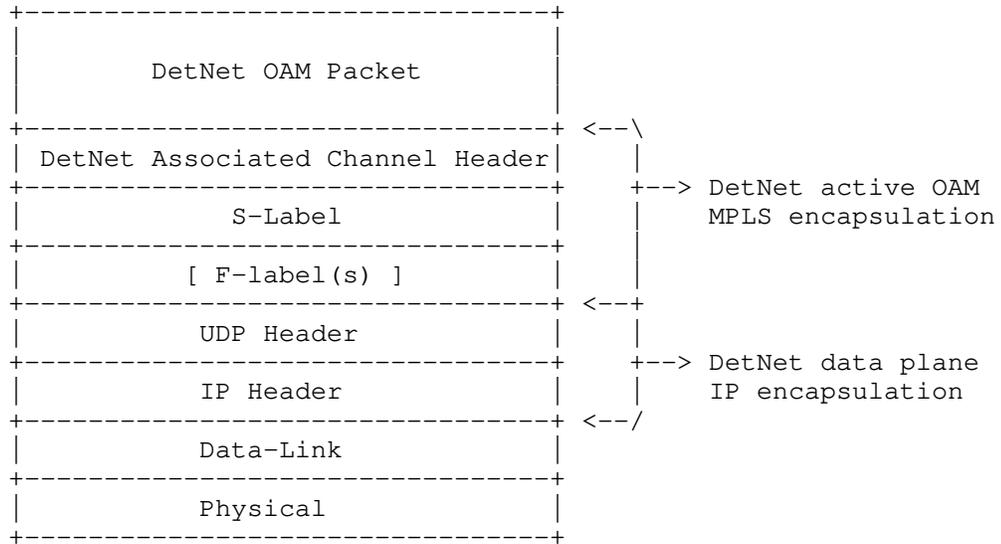


Figure 3: DetNet Active OAM Packet Encapsulation in MPLS-over-UDP/IP

Figure 4 displays the format of the DetNet Associated Channel Header (d-ACH).

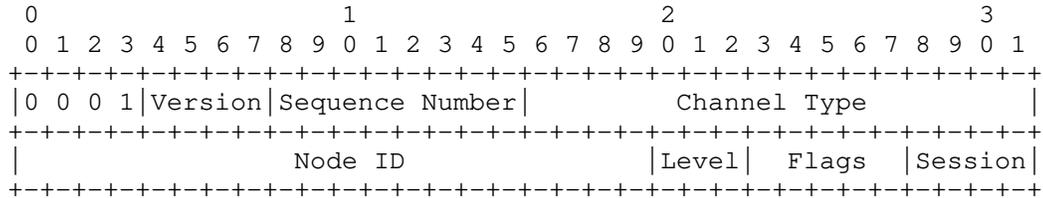


Figure 4: d-ACH Format

The d-ACH encodes the following fields:

Bits 0..3 MUST be 0b0001. This allows the packet to be distinguished from an IP packet [RFC4928] and from a DetNet data packet [RFC8964].

Version - a 4-bit field. This document specifies version 0.

Sequence Number - is an unsigned circular 8-bit field. Because a test packet of DetNet active OAM includes d-ACH, Section 4.2.1 of [RFC8964] does not apply to handling the Sequence Number field in DetNet OAM over the MPLS data plane. The sequence number space is circular with no restriction on the initial value. The originator DetNet node MUST set the value of the Sequence Number field before the transmission of a packet. the initial value SHOULD be random (unpredictable). The originator node SHOULD increase the value of the Sequence Number field by 1 for each active OAM packet. The originator MAY use other strategies, e.g., for negative testing of Packet Ordering Functions.

Channel Type - is a 16-bit field, and the value of DetNet Associated Channel Type. It MUST be one of the values listed in the IANA MPLS Generalized Associated Channel Types (including Pseudowire Associated Channel Types) registry [IANA-G-ACh-Types].

Node ID - is an unsigned 20-bit field. The value of the Node ID field identifies the DetNet node that originated the packet. A DetNet node MUST be provisioned with a Node ID that is unique in the DetNet domain. Methods of distributing Node ID are outside the scope of this specification.

Level - is a 3-bit field. Semantically, the Level field is analogous to the Maintenance Domain Level in [IEEE.802.1Q]. The Level field is used to cope with the "all active path forwarding" (defined by the TSN Task Group of the IEEE 802.1 WG [IEEE802.1TSNTG]) characteristics of the PREOF concept. A hierarchical relationship between OAM domains can be created using the Level field value, illustrated by Figure 18.7 in [IEEE.802.1Q].

Flags - is a 5-bit field. The Flags field contains five 1-bit flags. Section 5.1 creates the IANA d-ACH Flags registry for new flags to be defined. The flags defined in this specification are presented in Figure 5.

```

    0 1 2 3 4
    +--+--+--+
    |U|U|U|U|U|
    +--+--+--+
  
```

Figure 5: DetNet Associated Channel Header Flags Field Format

U: Unused and for future use. MUST be 0 on transmission and ignored on receipt.

Session ID is a 4-bit field. The Session field distinguishes OAM sessions originating from the same node (a given Maintenance End Point may have multiple simultaneously active OAM sessions) at the given Level.

A DetNet flow, according to [RFC8964], is identified by the S-Label that MUST be at the bottom of the stack. An Active OAM packet MUST include d-ACH immediately following the S-Label.

3.2. DetNet Packet Replication, Elimination, and Ordering Functions Interaction with Active OAM

At the DetNet service sub-layer, special functions (notably PREOF) MAY be applied to the particular DetNet flow to potentially reduce packet loss, improve the probability of on-time packet delivery, and ensure in-order packet delivery. PREOF relies on sequencing information in the DetNet service sub-layer. For a DetNet active OAM packet, PREOF MUST use the Sequence Number field value as the source of this sequencing information. App-flow and OAM use different sequence number spaces. PREOF algorithms are executed with respect to the sequence number space identified by the flow's characteristic information. Although the Sequence Number field in d-ACH has a range from 0 through 255, it provides sufficient space because the rate of

DetNet active OAM packet is significantly lower compared to the rate of DetNet packets in an App-flow; therefore, wrapping around is not an issue.

4. OAM Interworking Models

Interworking of two OAM domains that utilize different networking technology can be realized either by a peering or a tunneling model. In a peering model, OAM domains are within the corresponding network domain. When using the peering model, state changes that are detected by a Fault Management OAM protocol can be mapped from one OAM domain into another or a notification, e.g., an alarm, can be sent to a central controller. In the tunneling model of OAM interworking, usually, only one active OAM protocol is used. Its test packets are tunneled through another domain along with the data flow, thus ensuring the fate sharing among test and data packets.

4.1. OAM of DetNet MPLS Interworking with OAM of TSN

Active DetNet OAM can provide the end-to-end (E2E) fault management and performance monitoring for a DetNet flow. In the case of DetNet with an MPLS data plane and an IEEE 802.1 Time-Sensitive Networking (TSN) sub-network, this implies the interworking of DetNet active OAM with TSN OAM, which data plane aspects are specified in [RFC9037].

When the peering model (Section 4) is used in Connectivity Fault Management (CFM) OAM protocol [IEEE.802.1Q], then the node that borders both TSN and DetNet MPLS domains MUST support [RFC7023]. [RFC7023] specifies the mapping of defect states between Ethernet Attachment Circuits and associated Ethernet PWs that are part of an E2E emulated Ethernet service, and are also applicable to E2E OAM across DetNet MPLS and TSN domains. The CFM [IEEE.802.1Q] or in [ITU.Y1731] can provide fast detection of a failure in the TSN segment of the DetNet service. In the DetNet MPLS domain BFD (Bidirectional Forwarding Detection), specified in [RFC5880] and [RFC5885], can be used. To provide E2E failure detection, the TSN and DetNet MPLS segments could be treated as concatenated such that the diagnostic codes (see Section 6.8.17 of [RFC5880]) MAY be used to inform the upstream DetNet MPLS node of a failure of the TSN segment. Performance monitoring can be supported by [RFC6374] in the DetNet MPLS and [ITU.Y1731] in the TSN domains, respectively. Performance objectives for each domain should refer to metrics that is composable [RFC6049] or be defined for each domain separately.

The following considerations apply when using the tunneling model of OAM interworking between DetNet MPLS and TSN domains based on general principles described in Section 4 of [RFC9037]:

- * Active OAM test packets MUST be mapped to the same TSN Stream ID as the monitored DetNet flow.
- * Active OAM test packets MUST be treated in the TSN domain based on its S-Label and Class of Service marking (the Traffic Class field value).

Mapping between a DetNet flow and TSN Stream in the TSN sub-network is described in Section 4.1 of [RFC9037]. The mapping has to be done only on the edge node of the TSN sub-network, and intermediate TSN nodes do not need to recognize the S-Label. An edge node has two components:

1. A passive Stream identification function.
2. An active Stream identification function.

The first component identifies the DetNet flow (using Clause 6.8 of [IEEE.802.1CBdb]), and the second component creates the TSN Stream by manipulating the Ethernet header. That manipulation simplifies the identification of the TSN Stream in the intermediate TSN nodes by avoiding the need for them to look outside of the Ethernet header. DetNet MPLS OAM packets use the same S-Label as the DetNet flow data packets. The above-described mapping function treats these OAM packets as data packets of the DetNet flow. As a result, DetNet MPLS OAM packets are fate-sharing within the TSN sub-network. As an example of the mapping between DetNet MPLS and TSN, see Annex C.1 of [IEEE.802.1CBdb] that, in support of [RFC9037], describes how to match MPLS DetNet flows and TSN Streams can be achieved.

Note that the tunneling model of the OAM interworking requires that the remote peer of the E2E OAM domain supports the active OAM protocol selected on the ingress endpoint. For example, if BFD is used for proactive path continuity monitoring in the DetNet MPLS domain, BFD support (as defined in [RFC5885]) is necessary at any TSN endpoint of the DetNet service.

4.2. OAM of DetNet MPLS Interworking with OAM of DetNet IP

Interworking between active OAM segments in DetNet MPLS and DetNet IP domains can also be realized using either the peering or the tunneling model, as discussed in Section 4.1. Using the same protocol, e.g., BFD, over both segments, simplifies the mapping of errors in the peering model. For example, respective BFD sessions in DetNet MPLS and DetNet IP domains can be in a concatenated relationship as described in Section 6.8.17 of [RFC5880]. To provide performance monitoring over a DetNet IP domain, STAMP [RFC8762] and its extensions [RFC8972] can be used to measure packet loss and

packet delay metrics. Such performance metrics can be used to calculate composable metrics [RFC6049] within DetNet MPLS and DetNet IP domains to reflect the end-to-end DetNet service performance.

5. IANA Considerations

5.1. DetNet Associated Channel Header Flags Registry

This document describes a new IANA-managed registry to identify d-ACH Flags bits. The registration procedure is "IETF Review" [RFC8126]. The registry name is "DetNet Associated Channel Header (d-ACH) Flags". IANA should treat "DetNet Associated Channel Header (d-ACH) Flags" as the name of the registry group. There are five flags in the five-bit Flags field, defined as in Table 1.

Bit	Description	Reference
0-4	Unassigned	This document

Table 1: DetNet Associated Channel Header (d-ACH) Flags

6. Security Considerations

Security considerations discussed in DetNet specifications [RFC8655], [RFC9055], [RFC8964], and [I-D.ietf-detnet-oam-framework] are applicable to this document. Security concerns and issues related to MPLS OAM tools like LSP Ping [RFC8029], and BFD over PW [RFC5885] also apply to this specification.

7. Acknowledgment

Authors extend their appreciation to Pascal Thubert for his insightful comments and productive discussion that helped to improve the document. The authors are enormously grateful to Janos Farkas for his detailed comments and the inspiring discussion that made this document clearer and stronger. The authors recognize helpful reviews and suggestions from Andrew Malis, David Black, Tianran Zhou, and Kiran Makhijani. And special thanks are addressed to Ethan Grossman for his fantastic help in improving the document.

8. References

8.1. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7023] Mohan, D., Ed., Bitar, N., Ed., Sajassi, A., Ed., DeLord, S., Nigier, P., and R. Qiu, "MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking", RFC 7023, DOI 10.17487/RFC7023, October 2013, <<https://www.rfc-editor.org/info/rfc7023>>.
- [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>>.
- [RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.
- [RFC8964] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., Bryant, S., and J. Korhonen, "Deterministic Networking (DetNet) Data Plane: MPLS", RFC 8964, DOI 10.17487/RFC8964, January 2021, <<https://www.rfc-editor.org/info/rfc8964>>.
- [RFC9025] Varga, B., Ed., Farkas, J., Berger, L., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over UDP/IP", RFC 9025, DOI 10.17487/RFC9025, April 2021, <<https://www.rfc-editor.org/info/rfc9025>>.

8.2. Informational References

- [I-D.ietf-detnet-oam-framework]
Mirsky, G., Theoleyre, F., Papadopoulos, G. Z., Bernardos, C. J., Varga, B., and J. Farkas, "Framework of Operations, Administration and Maintenance (OAM) for Deterministic Networking (DetNet)", Work in Progress, Internet-Draft, draft-ietf-detnet-oam-framework-11, 8 January 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-detnet-oam-framework-11>>.
- [IANA-G-ACh-Types]
IANA, "MPLS Generalized Associated Channel (G-ACh) Types (including Pseudowire Associated Channel Types)", <<https://www.iana.org/assignments/g-ach-parameters/g-ach-parameters.xhtml#mpls-g-ach-types>>.

- [IEEE.802.1CBdb] IEEE, "IEEE Standard for Local and metropolitan area networks--Frame Replication and Elimination for Reliability Amendment 2: Extended Stream Identification Functions", IEEE 802.1CBdb, 2021.
- [IEEE.802.1Q] IEEE, "Bridges and Bridged Networks", IEEE 802.1Q, 2014.
- [IEEE802.1TSNTG] IEEE, "Time-Sensitive Networking (TSN) Task Group", IEEE 802.1Q, <<https://1.ieee802.org/tsn/>>.
- [ITU.Y1731] ITU-T, "OAM functions and mechanisms for Ethernet based Networks", ITU-T Recommendation G.8013/Y.1731, November 2013.
- [RFC3985] Bryant, S., Ed. and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, DOI 10.17487/RFC3985, March 2005, <<https://www.rfc-editor.org/info/rfc3985>>.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.
- [RFC4928] Swallow, G., Bryant, S., and L. Andersson, "Avoiding Equal Cost Multipath Treatment in MPLS Networks", BCP 128, RFC 4928, DOI 10.17487/RFC4928, June 2007, <<https://www.rfc-editor.org/info/rfc4928>>.
- [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>>.
- [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, <<https://www.rfc-editor.org/info/rfc5885>>.
- [RFC6049] Morton, A. and E. Stephan, "Spatial Composition of Metrics", RFC 6049, DOI 10.17487/RFC6049, January 2011, <<https://www.rfc-editor.org/info/rfc6049>>.

- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <<https://www.rfc-editor.org/info/rfc6374>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799, May 2016, <<https://www.rfc-editor.org/info/rfc7799>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8762] Mirsky, G., Jun, G., Nydell, H., and R. Foote, "Simple Two-Way Active Measurement Protocol", RFC 8762, DOI 10.17487/RFC8762, March 2020, <<https://www.rfc-editor.org/info/rfc8762>>.
- [RFC8972] Mirsky, G., Min, X., Nydell, H., Foote, R., Masputra, A., and E. Ruffini, "Simple Two-Way Active Measurement Protocol Optional Extensions", RFC 8972, DOI 10.17487/RFC8972, January 2021, <<https://www.rfc-editor.org/info/rfc8972>>.
- [RFC9037] Varga, B., Ed., Farkas, J., Malis, A., and S. Bryant, "Deterministic Networking (DetNet) Data Plane: MPLS over IEEE 802.1 Time-Sensitive Networking (TSN)", RFC 9037, DOI 10.17487/RFC9037, June 2021, <<https://www.rfc-editor.org/info/rfc9037>>.
- [RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <<https://www.rfc-editor.org/info/rfc9055>>.

Authors' Addresses

Greg Mirsky
Ericsson
Email: gregimirsky@gmail.com

Mach (Guoyi) Chen
Huawei
Email: mach.chen@huawei.com

Balazs Varga
Ericsson
Budapest
Magyar Tudosok krt. 11.
1117
Hungary
Email: balazs.a.varga@ericsson.com