

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 27 July 2026

X. Min
ZTE Corp.
G. Mirsky
Ericsson
L. Andersson
Bronze Dragon Consulting
23 January 2026

LSP Ping/Traceroute for Enabled In-situ OAM Capabilities
draft-ietf-mpls-lsp-ping-ioam-conf-state-01

Abstract

This document describes the application of the mechanism of discovering In-situ OAM (IOAM) capabilities, described in RFC 9359 "Echo Request/Reply for Enabled In Situ OAM (IOAM) Capabilities", in MPLS networks. The MPLS Node IOAM Information Query functionality uses the MPLS echo request/reply messages, allowing the IOAM encapsulating node to discover the enabled IOAM capabilities of each IOAM transit and IOAM decapsulating node.

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 27 July 2026.

Copyright Notice

Copyright (c) 2026 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
3. IOAM Capabilities Query TLV	3
3.1. Examples of the IOAM Capabilities Query	4
4. IOAM Capabilities Response TLV	4
4.1. IOAM Capabilities Sub-TLVs	5
4.2. Examples of IOAM Capabilities Response TLV	6
5. Return Code Field Processing	8
6. Operational Considerations	9
7. IANA Considerations	9
7.1. TLV assignments	9
7.2. New Sub-TLV registry	10
7.3. Return Code assignment	11
8. Security Considerations	11
9. Acknowledgements	12
10. References	12
10.1. Normative References	12
10.2. Informative References	13
Authors' Addresses	13

1. Introduction

MPLS encapsulation for In-situ OAM (IOAM) data is defined in [I-D.ietf-mpls-mna-ioam], which utilizes MPLS Network Actions (MNA) techniques ([RFC9789]) to carry IOAM data fields ([RFC9197], [RFC9326]) in MPLS packets.

As specified in [RFC9359], the echo request/reply can be used by the IOAM encapsulating node to discover the enabled IOAM capabilities at IOAM transit and decapsulating nodes.

[RFC8029] defines a probe message called "MPLS echo request", and a response message called "MPLS echo reply" for returning the result of the probe.

This document describes the MPLS Node IOAM Information Query functionality, which uses the MPLS echo request/reply messages, allowing the IOAM encapsulating node to discover the enabled IOAM capabilities of each IOAM transit and IOAM decapsulating node.

[RFC8029] specifies "ping" and "traceroute" modes. In "ping" mode, the ingress LSR sends a single MPLS echo request with the TTL in the outermost label stack entry set to 255. The MPLS echo request is intended to reach the end of the path and only the egress LSR is expected to respond with the MPLS echo reply. In "traceroute" mode, the ingress LSR transmits a sequence of MPLS echo requests with the TTL value being set in successive probe packets to 1, 2, and so on. Using TTL expiration as the exception mechanism, each LSR is expected to respond by transmitting an MPLS echo reply.

In an MPLS network, the ingress LSR may also act as the IOAM encapsulating node. In such a case, a transit LSR acts as the IOAM transit node, and the egress LSR acts as the IOAM decapsulating node. Usually, the trace option of IOAM data is needed, the IOAM encapsulating node requires to query the enabled IOAM capabilities of each IOAM transit and decapsulating node, then the "traceroute" mode can be used. In case that only the edge to edge option of IOAM data is needed, the IOAM encapsulating node requires to query the enabled IOAM Capabilities of only the IOAM decapsulating node, then the "ping" mode can be used.

The mechanism specified in this document applies to both point-to-point (P2P) MPLS LSP and point-to-multipoint (P2MP) MPLS LSP.

2. Conventions Used in This Document

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. IOAM Capabilities Query TLV

The IOAM Capabilities Query TLV presented in Figure 1 is carried as a TLV of the MPLS Echo Request message:

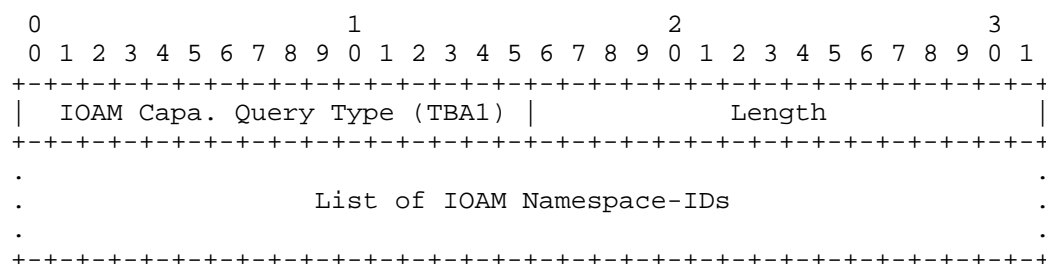


Figure 1: IOAM Capabilities Query TLV

Type: Indicates the IOAM Capabilities Query TLV. The value is TBA1.

Length: The length of the TLV's Value field in octets.

The Value field is a List of IOAM Namespace-IDs, which is also called IOAM Capabilities Query Container Payload in Section 3.1 of [RFC9359].

3.1. Examples of the IOAM Capabilities Query

The format of an IOAM Capabilities Query can vary from deployment to deployment.

In a deployment where only the default Namespace-ID is used, the IOAM Capabilities Query is depicted as the following:

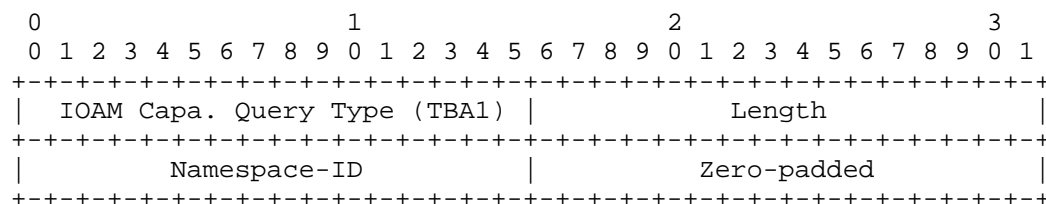


Figure 2: IOAM Capabilities Query of the Default IOAM Namespace

In a deployment where two Namespace-IDs (Namespace-ID1 and Namespace-ID2) are used, the IOAM Capabilities Query is depicted as the following:

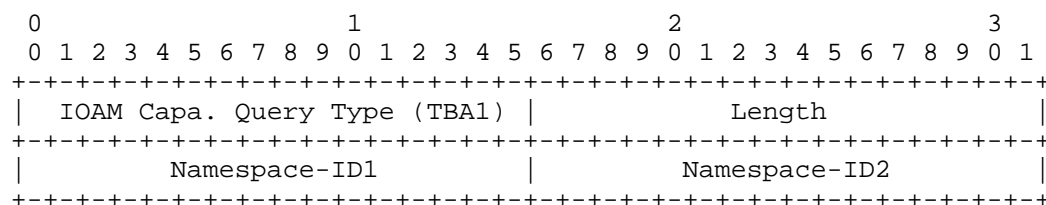


Figure 3: IOAM Capabilities Query of the Two IOAM Namespaces

4. IOAM Capabilities Response TLV

The IOAM Capabilities Response TLV presented in Figure 4 is carried as a TLV of the MPLS Echo Reply message:

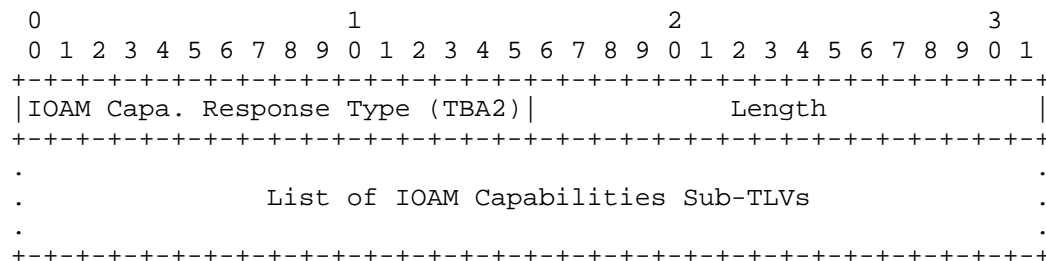


Figure 4: IOAM Capabilities Response TLV

Type: Indicates the IOAM Capabilities Response TLV. The value is TBA2.

Length: The length of the TLV's Value field in octets.

The Value field is a List of IOAM Capabilities Objects, which is also called IOAM Capabilities Response Container Payload in Section 3.2 of [RFC9359]. Each IOAM Capabilities Object is encoded in a sub-TLV format.

4.1. IOAM Capabilities Sub-TLVs

All IOAM Capabilities sub-TLVs (aka Objects) are encapsulated in an IOAM Capabilities Response TLV of an MPLS Echo Reply message.

Each IOAM Capabilities sub-TLV has the following format:

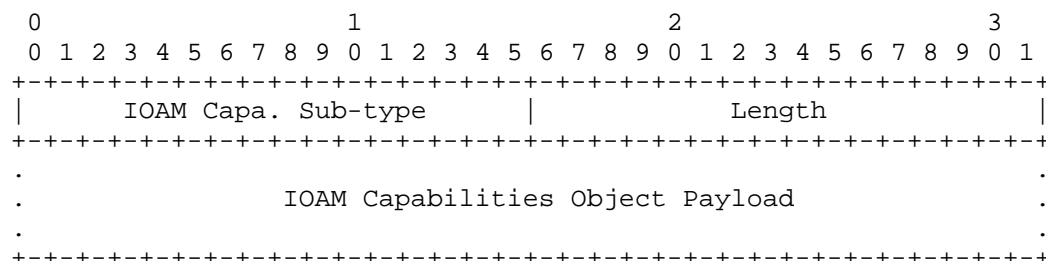


Figure 5: IOAM Capabilities Sub-TLV

Sub-type: Indicates the IOAM Capabilities sub-TLVs. The values are listed as the following:

Value -----	Sub-type Name -----
1	IOAM Pre-allocated Tracing Capabilities Object
2	IOAM Proof of Transit Capabilities Object
3	IOAM Edge-to-Edge Capabilities Object
4	IOAM DEX Capabilities Object
5	IOAM End-of-Domain Object

Length: The length of the sub-TLV's Value field in octets.

The Value field of the IOAM Capabilities sub-TLV is the IOAM Capabilities Object Payload, which is defined in Sections 3.2.1, 3.2.3, 3.2.4, 3.2.5, and 3.2.6 of [RFC9359].

4.2. Examples of IOAM Capabilities Response TLV

The format of an IOAM Capabilities Response can vary from deployment to deployment.

In a deployment where only the default Namespace-ID is used, the IOAM Pre-allocated Tracing Capabilities and IOAM Proof of Transit Capabilities are enabled at an IOAM transit node, if that IOAM transit node received an MPLS echo request containing IOAM Capabilities Query TLV, then the IOAM Capabilities Response TLV contained in an MPLS echo reply is depicted as the following:

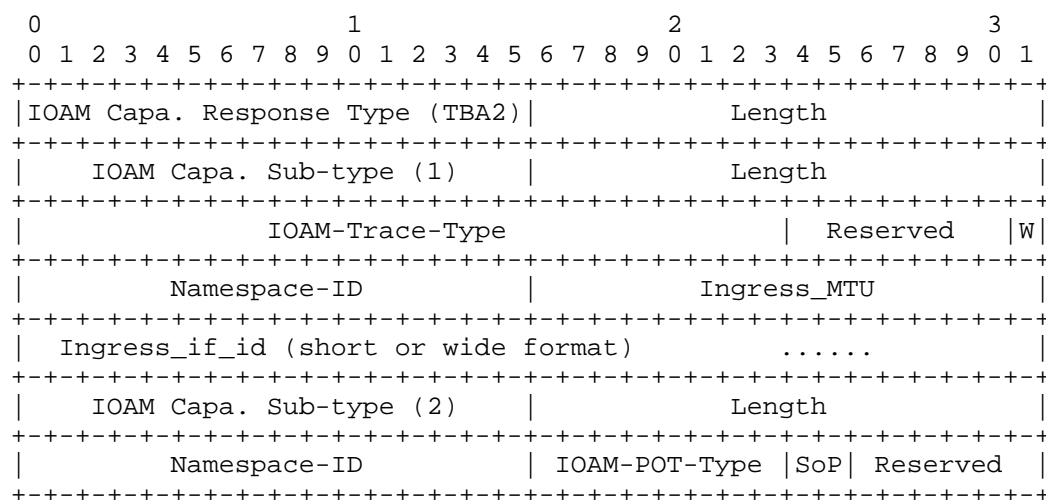


Figure 6: Example 1 of IOAM Capabilities Response TLV

In a deployment where two Namespace-IDs (Namespace-ID1 and Namespace-ID2) are used, for both Namespace-ID1 and Namespace-ID2 the IOAM Pre-allocated Tracing Capabilities and IOAM Proof of Transit Capabilities are enabled at an IOAM transit node, if that IOAM transit node received an MPLS echo request containing IOAM Capabilities Query TLV, then the IOAM Capabilities Response TLV contained in an MPLS echo reply is depicted as the following:

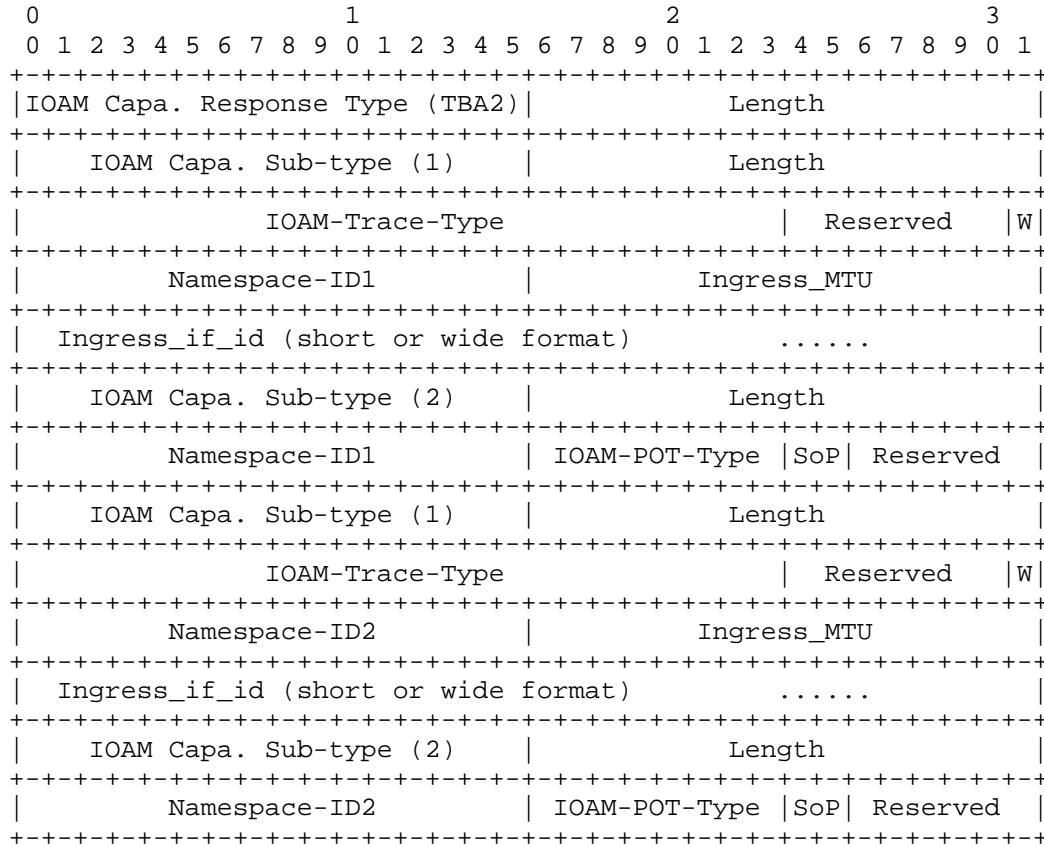


Figure 7: Example 2 of IOAM Capabilities Response TLV

Note that multiple sub-TLVs with the same sub-type may be present in an IOAM Capabilities Response TLV, as long as the Namespace-IDs in these sub-TLVs are all different.

In a deployment where only the default Namespace-ID is used, the IOAM Pre-allocated Tracing Capabilities, IOAM Proof of Transit Capabilities and IOAM Edge-to-Edge Capabilities are enabled at the IOAM decapsulating node, if that IOAM decapsulating node received an

MPLS echo request containing IOAM Capabilities Query TLV, then the IOAM Capabilities Response TLV contained in an MPLS echo reply is depicted as the following:

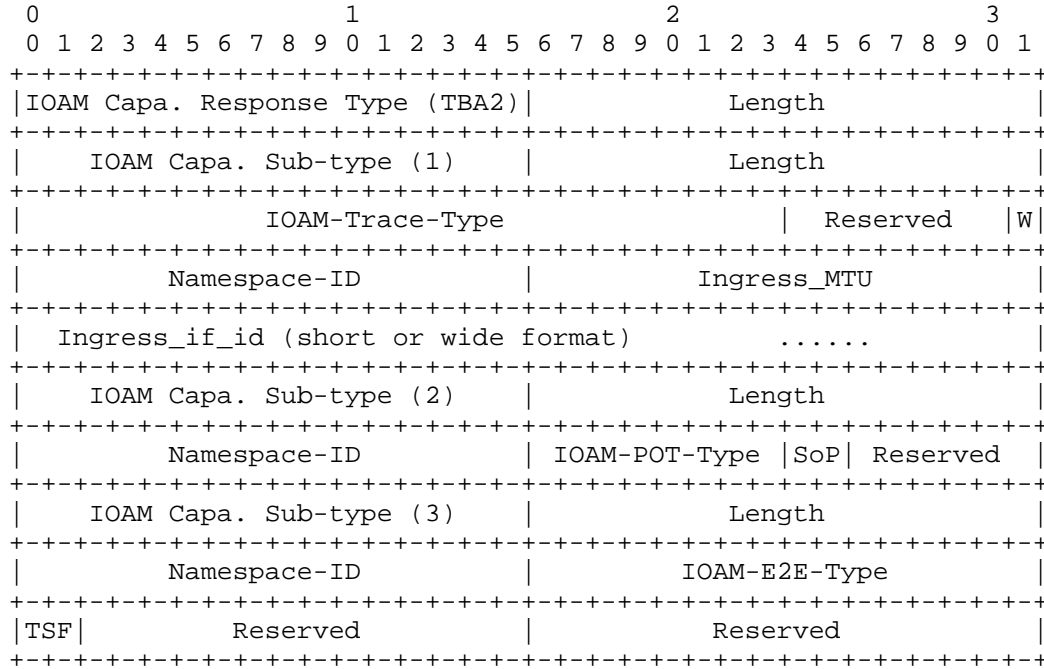


Figure 8: Example 3 of IOAM Capabilities Response TLV

5. Return Code Field Processing

The Return Code field in the MPLS echo reply MUST be set to (TBA3) No Matched Namespace-ID if any of the following conditions apply:

- * The IOAM Capabilities Query TLV does not include any Namespace-ID.
- * None of the contained list of IOAM Namespace-IDs is recognized.
- * None of the contained list of IOAM Namespace-IDs is enabled.

6. Operational Considerations

Section 4 of [RFC9359] provides an operational guide on how to use echo request/reply for discovering enabled IOAM capabilities at network nodes, which applies to this document as well. Specifically, if only the IOAM Edge-to-Edge Option-Type is enabled for an MPLS LSP, then LSP Ping in "ping" mode would be operated, which means the LSP ingress node would send a single MPLS echo request to the LSP egress node for discovering its IOAM Edge-to-Edge Capabilities; if other IOAM Option-Type than the Edge-to-Edge Option-Type is enabled for an MPLS LSP, e.g., the IOAM Pre-allocated Trace Option-Type is enabled for an MPLS LSP, then LSP Ping in "traceroute" mode would be operated, which means the LSP ingress node would send a sequence of MPLS echo requests with TTL equal to 1, 2, 3, and so on, until the LSP ingress node receives an MPLS echo reply sent by the LSP egress node.

Considering the IOAM function is resource-consuming, in an MPLS network with high number LSPs, usually not all the LSPs are IOAM enabled, in that case, the LSP Ping mechanism for IOAM capabilities discovery would be executed only with the IOAM-enabled LSPs.

7. IANA Considerations

This document does request that IANA assigns two new TLVs, and new sub-TLV registry for one of the new TLVs, 5 sub-TLVs to initially populate this registry and a new return code (TBA3).

7.1. TLV assignments

IANA is requested to assign two new TLVs (TBA1 and TBA2) from the "TLV" registry in the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry group. The TLVs' values should be assigned from the range of TLVs that require an error message if the TLV is not recognized. If possible the two lowest free values should be used for these TLVs.

Type	TLV Name	Reference	Sub-TLV Registry
TBA1	IOAM Capabilities Query	This document	No Sub-TLVs
TBA2	IOAM Capabilities Response	This document	Table 3

Table 1: New TLVs

7.2. New Sub-TLV registry

A new sub-TLV registry should be created for the TLV TBA2 created in Section 6.1.

The registration procedures for this sub-TLV registry shall be:

Range	Registration Procedure	Note
0-16383	Standards Action	This range is for TLVs that require an error message if not recognized. ([RFC9041], Section 3.1)
16384-31739	RFC Required	This range is for TLVs that require an error message if not recognized. ([RFC9041], Section 3.1)
31740-31743	Experimental Use	Reserved, not to be assigned. This range is for TLVs that require an error message if not recognized. ([RFC9041], Section 3.1)
31744-32767	First Come First Served	This range is for TLVs that require an error message if not recognized. ([RFC9041], Section 3.1)
32768-49161	Standards Action	This range is for TLVs that can be silently dropped if not recognized.
49162-64507	RFC Required	This range is for TLVs that can be silently dropped if not recognized.
64508-64511	Experimental Use	Reserved, not to be assigned. This range is for TLVs that can be silently dropped if not recognized.
64512-65535	First Come First Served	This range is for TLVs that can be silently dropped if not recognized.

Table 2: Sub-TLV Registration Procedures

This sub-TLV registry should initially be populated with the following values.

Sub-Type	Sub-TLV name	Reference	Comment
0	Reserved	This document	
1	IOAM Pre-allocated Tracing Capabilities Object	This document	
2	IOAM Proof of Transit Capabilities Object	This document	
3	IOAM Edge-to-Edge Capabilities Object	This document	
4	IOAM DEX Capabilities Object	This document	
5	IOAM End-of-Domain Object	This document	

Table 3: New Sub-TLV Registry for TLV TBA2

7.3. Return Code assignment

IANA is requested to assign a new Return Code from the "Return Code" registry in the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry group as follows:

Value	Meaning	Reference
TBA3	No Matched Namespace-ID	This document

Table 4: New Return Code

8. Security Considerations

Security issues discussed in [RFC8029] and [RFC9359] apply to this document.

This document recommends that the network operators establish policies that restrict access to MPLS Node IOAM Information Query functionality. In order to enforce these policies, nodes that support MPLS Node IOAM Information Query functionality are RECOMMENDED to support the following configuration options:

- * Enable/disable MPLS Node IOAM Information Query functionality. By default, MPLS Node IOAM Information Query functionality is disabled.
- * Define enabled Namespace-IDs. By default, all Namespace-IDs except the default one (i.e., Namespace-ID 0x0000) are disabled.

While applying the MPLS Node IOAM Information Query to P2MP MPLS LSP, since a single MPLS echo request may trigger multiple echo replies, there are scaling concerns and some mitigation measures, e.g., containing the Echo Jitter TLV in the MPLS echo request, as being specified in [RFC6425], MAY be applied.

9. Acknowledgements

The authors would like to acknowledge Tarek Saad for his comments on the idea of using LSP Ping for MPLS IOAM Capabilities Discovery.

The authors would like to acknowledge Adrian Farrel for his insightful review and comments.

10. References

10.1. Normative References

- [I-D.ietf-mppls-mna-ioam]
Gandhi, R., Mirsky, G., Li, T., Song, H., and B. Wen,
"Supporting In Situ Operations, Administration and
Maintenance Using MPLS Network Actions", Work in Progress,
Internet-Draft, draft-ietf-mppls-mna-ioam-04, 20 November
2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-mppls-mna-ioam-04>>.
- [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>>.
- [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>>.

- [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>>.
- [RFC9041] Andersson, L., Chen, M., Pignataro, C., and T. Saad, "Updating the MPLS Label Switched Paths (LSPs) Ping Parameters IANA Registry", RFC 9041, DOI 10.17487/RFC9041, July 2021, <<https://www.rfc-editor.org/info/rfc9041>>.
- [RFC9359] Min, X., Mirsky, G., and L. Bo, "Echo Request/Reply for Enabled In Situ OAM (IOAM) Capabilities", RFC 9359, DOI 10.17487/RFC9359, April 2023, <<https://www.rfc-editor.org/info/rfc9359>>.

10.2. Informative References

- [RFC9197] Brockners, F., Ed., Bhandari, S., Ed., and T. Mizrahi, Ed., "Data Fields for In Situ Operations, Administration, and Maintenance (IOAM)", RFC 9197, DOI 10.17487/RFC9197, May 2022, <<https://www.rfc-editor.org/info/rfc9197>>.
- [RFC9326] Song, H., Gafni, B., Brockners, F., Bhandari, S., and T. Mizrahi, "In Situ Operations, Administration, and Maintenance (IOAM) Direct Exporting", RFC 9326, DOI 10.17487/RFC9326, November 2022, <<https://www.rfc-editor.org/info/rfc9326>>.
- [RFC9789] Andersson, L., Bryant, S., Bocci, M., and T. Li, "MPLS Network Actions (MNAs) Framework", RFC 9789, DOI 10.17487/RFC9789, July 2025, <<https://www.rfc-editor.org/info/rfc9789>>.

Authors' Addresses

Xiao Min
ZTE Corp.
Nanjing
China
Phone: +86 18061680168
Email: xiao.min2@zte.com.cn

Greg Mirsky
Ericsson
United States of America
Email: gregimirsky@gmail.com

Loa Andersson
Bronze Dragon Consulting
Sweden
Email: loa@pi.nu