

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

L. Gong
China Mobile
C. Lin
New H3C Technologies
2 March 2026

Operations, Administration and Maintenance (OAM) for Network Resource
Partition (NRP) in MPLS Network
draft-gong-mpls-nrp-oam-mpls-02

Abstract

A Network Resource Partition (NRP) represents a subset of network resources and associated policies within the underlay network.

This document describes the implementation of the Operations, Administration, and Maintenance (OAM) mechanism for NRPs in MPLS networks. By extending existing OAM mechanisms such as ping, traceroute, the proposed solution enables comprehensive NRP support in MPLS networks.

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 3 September 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
1.1. Requirements Language	3
1.2. Terminology	3
2. OAM Mechanisms	3
3. MPLS PING	3
4. MPLS TRACEROUTE	5
5. UseCase	6
5.1. MPLS PING	6
5.2. MPLS TRACEROUTE	7
6. Security Consideration	8
7. IANA Considerations	8
8. Normative References	8
Authors' Addresses	8

1. Introduction

[RFC9543] provides the definition of IETF network slice for use within the IETF and discusses the general framework for requesting and operating IETF Network Slices, their characteristics, and the necessary system components and interfaces. It also introduces the concept Network Resource Partition (NRP), which is a subset of the resources and associated policies in the underlay network.

Using OAM tools enables real-time monitoring of the operational status of network slices, allowing for quick detection and localization of faults. When a node or link within a network slice experiences a failure, OAM tools can promptly issue alerts, assisting network administrators in taking swift corrective action to minimize service downtime. Therefore, the use of OAM tools in an NRP network is crucial for ensuring the availability and performance of network slice resources. This not only enhances user experience but also improves the overall efficiency and stability of the network.

Existing OAM tools typically include Ping, Traceroute. [RFC8029] describes how to Detect MPLS Data-Plane Failures in MPLS networks.

This document continues to employ these existing OAM mechanisms to monitor Data-Plane NRP resources Failures.

1.1. 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] (Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997) and [RFC8174] (Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017).

1.2. Terminology

The key terms used in this document are defined below.

Network Resource Partition (NRP): a subset of the network resources and associated policies on each of a connected set of links in the underlay network. This term is defined in [RFC9543].

IETF Network Slice: The realization of the service in the provider's network achieved by partitioning network resources and by applying certain tools and techniques within the network. This term is defined in [RFC9543].

2. OAM Mechanisms

[RFC8029] describes how to detect MPLS data-plane failures in MPLS networks.

To support NRP OAM, the initiator of an OAM operation (e.g., ping or traceroute) **MUST** include the NRP Identifier (NRP-ID) in the data plane of the probe packets. The method for carrying the NRP-ID in the data plane is outside the scope of this document; it is assumed that the underlay network provides a mechanism to associate a packet with a specific NRP (e.g., through a dedicated label, a traffic class, or a slice identifier).

Intermediate equipment and OAM End Points need to check the NRP resources when receiving OAM packets with an NRP-ID. If the resource check fails, the node shall respond with an MPLS Echo Reply carrying an appropriate error code (see Section 7)

3. MPLS PING

When performing an MPLS PING operation, the initiator sends an MPLS Echo request. To support probing of NRP resources, the NRP-ID **MUST** be carried in the data plane (e.g., via a dedicated label or other encoding). The MPLS Echo request is forwarded along the LSP towards the destination.

The intermediate node processes the MPLS Echo Request, looks up the forwarding table, obtains the Downstream information, and checks the NRP to the Downstream. If NRP are not available, it responds with a MPLS Echo Reply, indicating the Error as "NRP resources unavailable". If this node does not recognize the NRP ID or has not allocated resources for this NRP, it returns "NRP unknown or not supported".

According to [RFC8029], an MPLS Echo Request is designed to be processed in the control plane of transit nodes and the egress node, triggered by mechanisms such as the Router Alert Option, the MPLS Router Alert label, or TTL expiration. This document does not alter that fundamental behavior. The newly added error codes apply only to scenarios where the packet is processed in the control plane.

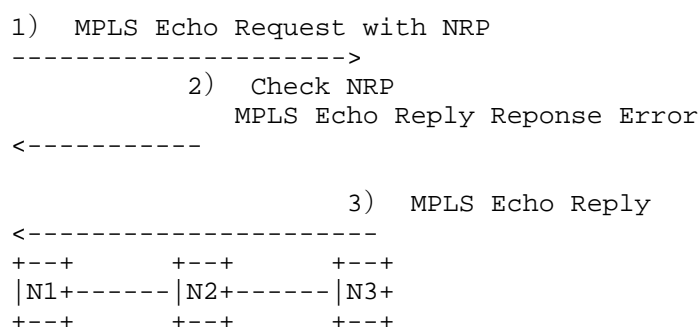


Figure 1: MPLS PING for NRP

Process of MPLS PING for NRP:

1) The initiator of the MPLS Echo Request includes the NRP-ID in the data layer when sending the MPLS PING request.

2) The intermediate node processes the MPLS Echo Request, looks up the forwarding table, obtains the Downstream information, and checks the NRP to the Downstream. If NRP are not available, it responds with a MPLS Echo Reply, indicating the Error as "NRP resources unavailable". If this node does not recognize the NRP ID or has not allocated resources for this NRP, it returns "NRP unknown or not supported".

For MPLS networks, it is necessary to extend the Return Codes carried in the MPLS Echo Reply(IANA 7).

3) If the check passes, the End Point will respond with a normal MPLS Echo Reply.

4. MPLS TRACEROUTE

When performing a MPLS TRACEROUTE operation, the TRACEROUTE initiator sends MPLS Echo request packets toward the destination node by incrementally increasing the TTL value. To support the probing of NRP resources, NRP information is carried in the data layer. Each intermediate node, when forwarding the MPLS Echo request, looks up the forwarding table to obtain the Downstream information and performs NRP check for the next hop. If the resources are unavailable, the node responds with an MPLS Echo Reply with error message indicating "NRP resource unavailability". If this node does not recognize the NRP ID or has not allocated resources for this NRP, it returns "NRP unknown or not supported". The packets used for MPLS TRACEROUTE are the same as those used for MPLS PING. When NRP resources are unavailable, the error codes used are also identical to those used in MPLS PING operations

```

1)          MPLS Echo Request with NRP-ID
----->
      2)    MPLS Echo Reply
<-----
3)    MPLS Echo Request with NRP-ID
----->
          4)    MPLS Echo Reply
<-----
+---+      +---+      +---+
|N1+-----|N2+-----|N3+
+---+      +---+      +---+

```

Figure 2: MPLS Traceroute for NRP

Process of MPLS Traceroute for NRP:

1) The initiator of the MPLS Echo request includes the NRP-ID in the data layer when sending the Traceroute request.

The MPLS Echo Request with TTL 1 to n increase.

2) The intermediate node looks up the forwarding table to obtain the Downstream information and performs NRP check for the Downstream when processing a MPLS Echo Request. If they are not available, it responds with a MPLS Echo Reply, indicating the Error as "NRP resources unavailable". If this node does not recognize the NRP ID or has not allocated resources for this NRP, it returns "NRP unknown or not supported". The error code for expansion should be the same as MPLS PING.

3) If the check passes, the process proceeds with a normal MPLS Traceroute, performing hop-by-hop detection of the path to the End Point until the Traceroute process is completed, and the detection results are outputted.

5. UseCase

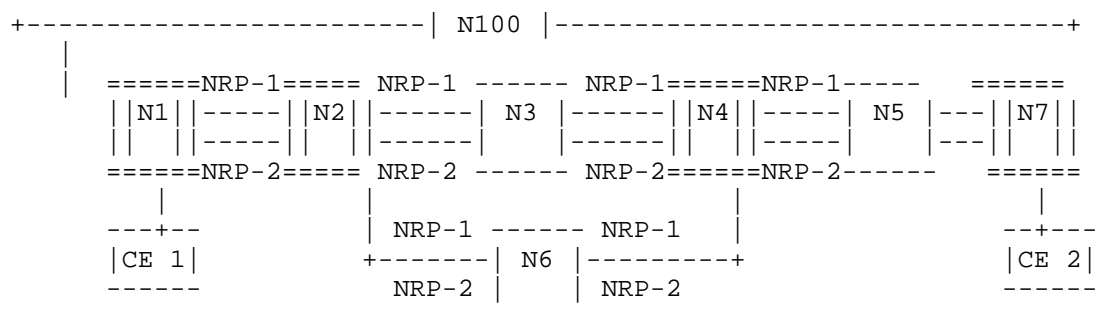


Figure 3: NRP Network Diagram

In the reference topology of Figure 3:

- * Node j has an IPv4 loopback address 192.168.j.1/32.
- * A label at node j is 1j000 (e.g., node 2 uses label 12000).
- * Node N100 is a controller.

Two NRPs are defined: - NRP-1 (ID=1) and NRP-2 (ID=2). Different links and nodes may participate in different NRPs as shown.

The following subsections illustrate MPLS ping and traceroute operations with NRP support.

5.1. MPLS PING

Example 1: Successful MPLS ping through NRP-1.

```
ping 15000 via label-stack 12000, 14000, NRP-ID: 1, Ret NRP-ID: 2
```

Sending 5, 100-byte MPLS Echos to 192.168.5.1, timeout is 2 seconds:

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 0.625
/0.749/0.931 ms
```

Example 2: MPLS ping failure due to NRP resource unavailability.

ping 15000 via label-stack 12000, 14000, NRP-ID: 1, Ret NRP-ID: 2

Reply to request 2 (1 ms) from 192.168.2.1. Return Code: 'NA'

Reply to request 3 (1 ms) from 192.168.2.1. Return Code: 'NA'

Reply to request 4 (1 ms) from 192.168.2.1. Return Code: 'NA'

Reply to request 5 (1 ms) from 192.168.2.1. Return Code: 'NA'

Success rate is 0 percent (0/5), round-trip min/avg/max = 1/1/1 ms

In the above examples: - 'NA' indicates "NRP resources unavailable". These codes are placeholders for the IANA-assigned values (see Section 7).

5.2. MPLS TRACEROUTE

Example 1: Successful MPLS traceroute through NRP-1.

traceroute 15000 via label-stack 12000, 14000, NRP-ID: 1, Ret-NRP-ID: 2

Tracing the route to 15000

TTL	Replier	Time	Type	Downstream
0			Ingress	192.168.2.1/[12000]
1	192.168.2.1	1 ms	Transit	192.168.4.1/[14000]
2	192.168.4.1	1 ms	Transit	192.168.5.1/[15000]
3	192.168.5.1	1 ms	Egress	

Example 2: MPLS traceroute failure due to NRP resource unavailability.

> traceroute 15000 via label-stack 12000, 14000, NRP-ID: 1, Ret-NRP-ID: 2

Tracing the route to 15000

TTL	Replier	Time	Type	Downstream
0			Ingress	192.168.2.1/[12000]
1	192.168.2.1	1 ms	Transit	192.168.4.1/[14000] Return[NA]

6. Security Consideration

This document does not impose any additional security challenges beyond those described in [RFC4884], [RFC4443], [RFC0792], [RFC8754], and [RFC8986]. The inclusion of an NRP-ID in OAM packets does not introduce new vulnerabilities, as the NRP-ID is used only within the trusted domain of a network provider. Operators should ensure that OAM packets are adequately protected (e.g., by filtering at network boundaries) to prevent unauthorized injection or disclosure of network slice information.

7. IANA Considerations

IANA is requested to allocate two new Return Codes in the "Multi-Protocol Label Switching (MPLS) Label Switched Path (LSP) Ping Parameters" registry, sub-registry "Return Codes".

The following values are proposed (specific code points to be assigned by IANA):

Value	Meaning
-----	-----
TBD1	NRP unknown/not supported
TBD2	NRP resource unavailable

The code "NRP unknown/not supported" indicates that the egress LSR does not recognize the NRP-ID or the NRP is not provisioned on that node. The code "NRP resource unavailable" indicates that the NRP is recognized but the required resources (e.g., bandwidth, queue) are not currently available.

8. Normative References

- [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/rfc/rfc8029>>.
- [RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/rfc/rfc9543>>.

Authors' Addresses

Liyan Gong
China Mobile
China
Email: gongliyan@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com