

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

Z. Li
J. Dong
Huawei Technologies
2 March 2026

Carrying NRP related Information in MPLS Packets
draft-ietf-mpls-mna-psd-nrp-selector-01

Abstract

A Network Resource Partition (NRP) is a subset of the network resources and associated policies on each of a connected set of links in the underlay network. An NRP could be used as the underlay to support one or a group of enhanced VPN services. Multiple NRPs can be created by network operator to meet the diverse requirements of different enhanced VPN services. In packet forwarding, some fields in the data packet needs to be used as the NRP Selector to identify the NRP the packet belongs to, so that the NRP-specific processing can be executed on the packet.

This document proposes a mechanism to carry the NRP Selector ID and related information in MPLS packets. The procedure for processing the NRP Selector ID and related information is also specified.

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
2. Carrying NRP Information in MPLS Packets	3
3. Procedures	5
3.1. PS-NRP Action Insertion	5
3.2. NRP based Packet Forwarding	5
4. Capability Advertisement and Negotiation	6
5. IANA Considerations	6
6. Security Considerations	6
7. Contributors	7
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	8
Authors' Addresses	9

1. Introduction

Virtual Private Networks (VPNs) [RFC4206] provide different customers with logically isolated connectivity over a common network infrastructure. With the introduction of 5G and also in some existing network scenarios, some customers may require network connectivity services with advanced features comparing to conventional VPNs, such as resource isolation from other services or guaranteed performance. Such kind of network service is called enhanced VPN [RFC9732]. Enhanced VPN service requires the coordination and integration between the overlay VPNs and the capability and resources of the underlay network. Enhanced VPN can be used, for example, to deliver network slice services as described in [RFC9543].

[RFC9543] also introduces the concept of the Network Resource Partition (NRP), which is a subset of the buffer/queuing/scheduling resources and associated policies on each of a connected set of links in the underlay network. An NRP can be associated with a logical network topology to select or specify the set of links and nodes involved.

In packet forwarding, traffic of different Enhanced VPN services needs to be processed separately based on the network resources and the logical topology associated with the corresponding NRP. [I-D.ietf-teas-nrp-scalability] describes the scalability considerations and the possible optimizations for providing a relatively large number of NRPs. The approach to improve the data plane scalability of NRP is to introduce a dedicated data plane identifier (which is called NRP Selector ID) in the data packets to identify the set of network resources allocated to an NRP, so that the packets mapped to an NRP can be processed and forwarded using the NRP-specific network resources, which could avoid possible resource competition with other services in the same network.

This document proposes a mechanism to carry the NRP Selector ID and related information in MPLS [RFC3031] data packets, so that the packet will be processed by network nodes using the subset of network resources and policies of the corresponding NRP. The procedure for processing the NRP Selector ID is also specified. The destination and forwarding path of the MPLS LSP is determined using the MPLS label stack in the packet, and the set of network resources and policies used for processing the packet is determined by the NRP action carried as post-stack data (PSD) [I-D.ietf-mpls-mna-ps-hdr]. The mechanism introduced in this document is applicable to both MPLS networks with RSVP-TE [RFC3209] or LDP [RFC5036] LSPs, and MPLS networks with Segment Routing (SR) [RFC8402] [RFC8660].

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 BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Carrying NRP Information in MPLS Packets

This document makes use of the post-stack header mechanism as defined in [I-D.ietf-mpls-mna-ps-hdr]. A new type of Post-Stack network action called "Network Resource Partition (NRP)" is defined to carry the NRP Selector ID and NRP related information. The Opcode of the NRP action is to be assigned by IANA. The format of the NRP action and its ancillary data is shown as below:

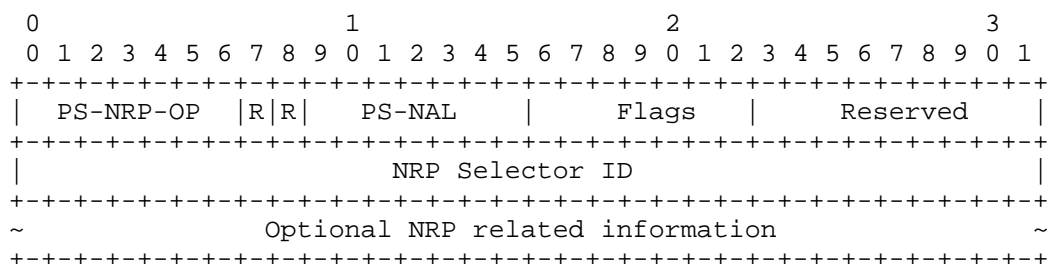



Figure 1. The format of MPLS NRP Action

Where:

PS-NRP-OP: Post-stack Opcode for the NRP action, the value is to be assigned by IANA.

PS-NAL: Post-Stack network action length in 4-octet units, excluding the first 4-octets.

Flags: 8-bit flag field. The most significant bit is defined in this document:



- S (Strict Match): The S flag is used to indicate whether the NRP Selector ID MUST be strictly matched for the processing of the packet. When the S flag in the NRP Action of a received packet is set to 1, if the NRP Selector ID in the packet does not match with any of the network resources provisioned on the network node, the packet MUST be dropped. When the S flag in the NRP Action of a received packet is set to 0, if the NRP Selector ID in the packet does not match with any of the network resources provisioned on the network node, the packet MUST be forwarded using the default set of resource and behavior as if the NRP Action does not exist.
- U (Unused): These flags are reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

Reserved: 8-bit field reserved for future use.

NRP Selector ID: 4-octet network-wide identifier which uniquely identifies an NRP.

Optional NRP related information: NRP related information which may be carried as ancillary data of the NRP action. One example is the traffic accounting data of a specific NRP. The detailed encoding and mechanisms are out of the scope of this document. The length of the optional NRP related information in 4-octet units is determined by the value of PS-NAL minus 1.

The PS-NRP action SHOULD be processed hop-by-hop (HBH). It is suggested the PS-NRP action be placed closer to the top of the PSD than any Post-Stack actions with the Ingress-To-Egress scope.

The NRP Selector ID and related information are carried as Ancillary Data of Post-Stack action due to the following reasons:

- * NRP Selector ID with 32-bit length does not fit well into the MPLS Label Stack Entry (LSE) due to the existence of S bit in LSE (e.g., in 31-bit Format D).
- * Optional NRP related information which may be mutable or variable is more suitable and efficient to be encoded as Post-stack Data.

3. Procedures

3.1. PS-NRP Action Insertion

When an LSP ingress receives a packet, the traffic classifications and mapping policies are checked. If a match is found that the packet needs to be steered on to one of the NRPs of the MPLS network, the packet is encapsulated with an MPLS Label Stack and a Post-Stack Header [I-D.ietf-mpls-mna-ps-hdr] is added after the label stack. The Post-Stack Header contains a Post-Stack NRP action, which includes the NRP Selector ID and optional NRP related information.

3.2. NRP based Packet Forwarding

On receipt of an MPLS packet which carries the PS-NRP action, network nodes which support the mechanism defined in this document MUST parse the Post-Stack header and the NRP action, and use the NRP Selector ID to identify the NRP the packet belongs to, then the local resources allocated to the NRP are used to process and forward the packet. The forwarding behavior is based on both the MPLS label stack and the Post-Stack NRP action. The top MPLS label is used for the lookup of the next-hop, and the NRP Selector ID can be used to determine the set of network resources allocated by the network nodes for processing and sending the packet to the next-hop. NRP-specific information such as NRP based packet statistics may be carried as optional ancillary data in the NRP action, the details are out of the

scope of this document. Network nodes which do not support the mechanism in this document MUST ignore the NRP action, and forward the packet only based on the top MPLS label.

There can be different approaches used for allocating network resources on each network node to the NRPs. For example, on one interface, a subset of forwarding plane resources (e.g., bandwidth and the associated buffer/queuing/scheduling resources) allocated to a particular NRP can be considered as a virtual layer-2 sub-interface with dedicated bandwidth and the associated resources. In packet forwarding, the top MPLS label of the received packet is used to identify the next-hop and the outgoing layer-3 interface, and the NRP Selector ID is used to further identify the virtual sub-interface which is associated with the NRP on the outgoing interface.

The egress node of the MPLS LSP MUST pop the Post-Stack NRP action, together with the Post-Stack header and other Post-Stack actions if there is any.

4. Capability Advertisement and Negotiation

Before inserting the Post-Stack NRP action into an MPLS packet, the ingress node wants to know whether the nodes along the LSP can process the NRP extension header properly based on the mechanisms defined in this document. This can be achieved by introducing the capability advertisement and negotiation mechanism for the Post-Stack NRP action. The ingress node also needs to know whether the egress node of the LSP can remove the Post-Stack NRP action as part of the Post-Stack header properly before parsing the upper layer and sending the packet to the next hop. The signaling mechanism for capability advertisement and negotiation is out of the scope of this document.

5. IANA Considerations

IANA is requested to assign the code point for the PS-NRP action from the "Post-Stack Network Action Opcodes" registry as below:

Value	Description	Reference

TBA	Post-Stack Network Action for NRP	this document

6. Security Considerations

The security considerations in [RFC3032] and [I-D.ietf-mpls-mna-ps-hdr] apply to this document.

The introduction of Post-Stack NRP actions allows attacking traffic targeting at a specific NRP in the network, which can impact the performance of services carried by the NRP. Operator needs to make sure that only trusted network nodes can add Post-stack NRP action to packets.

7. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

8. Acknowledgements

The authors would like to thank Loa Andersson for the review and comments.

9. References

9.1. Normative References

- [I-D.ietf-mpls-mna-ps-hdr]
Rajamanickam, J., Gandhi, R., Zigler, R., and J. Dong,
"Post-Stack MPLS Network Action (MNA) Solution", Work in
Progress, Internet-Draft, draft-ietf-mpls-mna-ps-hdr-06,
20 January 2026, <[https://datatracker.ietf.org/doc/html/
draft-ietf-mpls-mna-ps-hdr-06](https://datatracker.ietf.org/doc/html/draft-ietf-mpls-mna-ps-hdr-06)>.
- [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>>.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol
Label Switching Architecture", RFC 3031,
DOI 10.17487/RFC3031, January 2001,
<<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y.,
Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack
Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001,
<<https://www.rfc-editor.org/info/rfc3032>>.
- [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>>.

- [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/info/rfc9543>>.
- [RFC9732] Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for NRP-Based Enhanced Virtual Private Networks", RFC 9732, DOI 10.17487/RFC9732, March 2025, <<https://www.rfc-editor.org/info/rfc9732>>.

9.2. Informative References

- [I-D.ietf-teas-nrp-scalability]
Dong, J., Li, Z., Gong, L., Yang, G., and G. S. Mishra, "Scalability Considerations for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-scalability-09, 11 February 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-09>>.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, DOI 10.17487/RFC4206, October 2005, <<https://www.rfc-editor.org/info/rfc4206>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: robinli314@163.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Road
Beijing
100095
China
Email: jie.dong@huawei.com