Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

J. Dong
Z. Li
Huawei Technologies
C. Xie
C. Ma
China Telecom
G. Mishra
Verizon Inc.

7 July 2025

Carrying Network Resource (NR) related Information in IPv6 Extension Header
draft-ietf-6man-enhanced-vpn-vtn-id-12

Abstract

Virtual Private Networks (VPNs) provide different customers with logically separated 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 VPN services. Such kind of network service is called enhanced VPNs. Enhanced VPNs can be used, for example, to deliver network slice services.

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 may be used as the underlay to support one or a group of enhanced VPN services. For packet forwarding within a specific NRP, some fields in the data packet need to be used to identify the NRP to which the packet belongs. In doing so, NRP-specific processing can be performed on each node along the forwarding path in the NRP.

This document specifies a new IPv6 Hop-by-Hop option to carry network resource related information (e.g., identifier) in data packets. The NR Option can be used to carry NRP Selector ID and related information, while it is designed to make the NR option generalized for other network resource semantics and functions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

This Internet-Draft will expire on 8 January 2026.

Copyright Notice

Table of Contents

1. Introduction

Virtual Private Networks (VPNs) [RFC4026] 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]. The realization of enhanced VPN services require the coordination and integration between the overlay VPNs and the capability and resources of the underlay network. Enhanced VPNs can be used, for example, to deliver Network Slice Services as described in Section 7.4 of [RFC9543].

Section 7.1 of [RFC9543] 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 may be associated with a logical network topology to select or specify the set of links and nodes involved.

[RFC9732] specifies the framework of NRP-based enhanced VPN and describes the candidate component technologies in different network planes and network layers. An NRP could be used as the underlay to meet the requirement of one or a group of enhanced VPN services.

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. One approach to improve the data plane scalability of NRPs is to introduce a dedicated NRP Selector ID in data packets, which is used to identify the set of network resources allocated to an NRP. This way, packets mapped to an NRP can be processed and forwarded using the NRP-specific network resources, which could help to provide guaranteed performance for the packets. An NRP Selector ID can be used to identify a subset of the resources (e.g., bandwidth, buffer, and queuing resources) allocated on the set of links and nodes involved in the NRP. The logical topology associated with an NRP could be defined and identified using mechanisms such as Multi-Topology [RFC4915], [RFC5120], or Flex-Algo [RFC9350].

This document specifies a mechanism to carry network resource related information in a new IPv6 Hop-by-Hop option (Section 4.3 of [RFC8200]) called "Network Resource (NR) option". In networks built

with NRPs, the NR option SHOULD be parsed by every intermediate node along the forwarding path, and the obtained NRP Selector ID is used to invoke NRP-specific packet processing and forwarding using the set of NRP-specific resources. This solution is designed to support a large number of NRPs in IPv6 networks [I-D.ietf-teas-nrp-scalability].

In this document the application of the NR option is to indicate the NRP-specific resource information, while the NR option is considered as a generic mechanism to convey network-wide resource ID and information with different semantics and functions to meet the possible use cases in the future. Some considerations about the generalization of the NR Option are described in Section 5.

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. IPv6 Extension Header Option for Network Resource Information

A new Hop-by-Hop option (Section 4.3 of [RFC8200]) type "Network Resource" is defined to carry the network resource related information. Its format is shown in Figure 1.

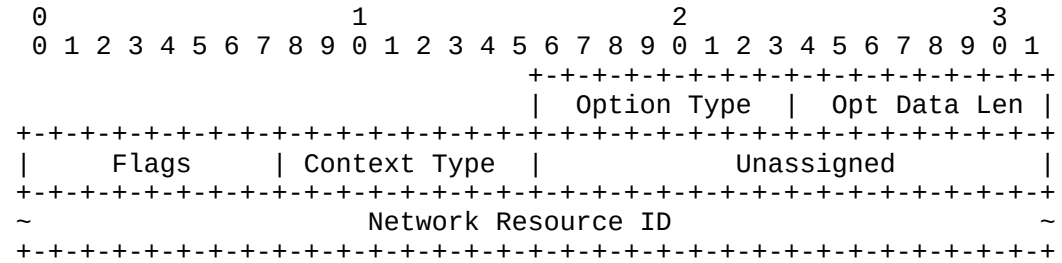

Figure 1. The format of Network Resource (NR) Option

Option Type: 8-bit identifier of the type of option. The type of NR option is TBA. The bits of the type field are defined as shown below:

* BB 00 The highest-order 2 bits are set to 00 to indicate that a node which does not recognize this type will skip over it and continue processing the header.

* C 0 The third highest-order bit is set to 0 to indicate this option does not change en route.

* TTTTT tba.

Opt Data Len: 8-bit unsigned integer indicates the length of the option Data field of this option, in octets.

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

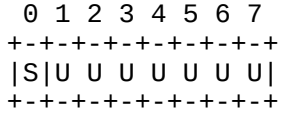

* S (Strict Match): The S flag is used to indicate whether the NR ID MUST be strictly matched for the processing of the packet. When the S flag in the NR option of a received packet is set to 1, if the NR ID in the packet does not match with any of the network resources provisioned on the network node, the packet MUST be dropped. When the S flag in the NR option of a received packet is set to 0, if the NR 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 NR option does not exist.

* U (Unassigned): These flags are reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

The setting of the S flag depends on the operator’s policy. Such policy can be NRP-specific, and may be at a fine granularity to apply to a subset of packets within an NRP. Such policy needs to be provided to the ingress nodes to apply to packets which are mapped to corresponding NRPs. For a given NRP, the suggested default policy is to make the S flag set.

As an example, for OAM packets which are used to detect the availability of a forwarding path associated with NRP-specific resources, the S flag should be set to 1. This way, only when the set of network resources and policies are correctly instantiated for the NRP on all network links along a path, the OAM packets can be received by an egress endpoint and the availability check can be passed.

The S flag in the NR option provides an approach for fine-granular control of the forwarding policy of packets whose NR IDs do not match with the network resources provisioned on the transit network nodes.

One alternate approach is to specify the forwarding policy of packets in different NRPs via configuration, while additional configuration would be needed when non-default fine-granular policy is required for a given NRP.

Context Type (CT): One-octet field used to indicate the semantics of the NR ID carried in the option. The context value defined in this document is as follows:

* CT=0: The NR ID is a network-wide unique NRP Selector ID, which is used to identify the subset of network resources allocated to the NRP on the involved network links.

Unassigned: 2-octet field reserved for future use. They MUST be set to 0 on transmission and MUST be ignored on receipt.

NR ID: The identifier of a set of network resources, the semantics of the ID is determined by the Context Type. The length of the NR ID is the Opt Data Length minus 4.

Note that, in the context of 5G network slicing, if a deployment found it useful, a four-octet NRP Selector ID field (CT=0) may be derived from the four-octet Single Network Slice Selection Assistance Information (S-NSSAI) defined in 3GPP [TS23501].

3. Procedures

This section describes the procedures for NR option processing when the value of the Context Type (CT) is set to 0. In this case the NRP Selector ID is carried in the NR option. The processing procedures for NR option with other CT values are out of the scope of this document, and will be specified in separate documents which introduce those CT values.

3.1. Adding NR Option to Packets

When an ingress node of an IPv6 domain receives a packet, according to the traffic classification and mapping policy, if the packet needs to be steered into an NRP, then the packet MUST be encapsulated in an outer IPv6 header with the source and destination addresses set according to the local policy. The NRP Selector ID which the packet is mapped to according to the policy MUST be carried in the NR option of the Hop-by-Hop Options header, which is associated with the outer IPv6 header. It is RECOMMENDED the NR option is carried as the first option in the Hop-by-Hop Options header.

3.2. NRP-specific Packet Forwarding

On receipt of a packet with an NR option, each transit network nodes which can process the Hop-by-Hop Options header and the NR option at full forwarding rate [RFC9673] MUST use the NRP Selector ID to determine the set of local network resources which are allocated to the NRP. The packet forwarding behavior is based on both the destination IP address and the NRP Selector ID. More specifically, the destination IP address SHOULD be used to determine the next-hop and the outgoing interface, and the NRP Selector ID SHOULD be used to determine the subset of network resources on the outgoing interface which are allocated to the NRP for processing and sending the packet. If the NRP Selector ID in the packet does not match with any of the NRP provisioned on the outgoing interface, the S flag in the NR option SHOULD be used to determine whether the packet should be dropped or forwarded using the default set of network resources of the outgoing interface. The Traffic Class field of the outer IPv6 header MAY be used to provide differentiated treatment for packets which belong to the same NRP. On the egress nodes of the IPv6 domain, if the destination address in the outer IPv6 header of a received packet matches with a local address, it MUST decapsulate the outer IPv6 header and the Hop-by-Hop Options header which includes the NR option.

There can be different approaches of partitioning the network resources and allocating them to different NRPs in the forwarding plane. For example, on one physical interface, a subset of the forwarding plane resources (e.g., bandwidth and the associated buffer and queuing resources) can be allocated to a particular NRP and represented as a virtual sub-interface or a data channel with reserved bandwidth resource. The IPv6 destination address 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 or the data channel on the outgoing interface which is associated with the NRP.

According to [RFC9673], network nodes which do not support the processing of Hop-by-Hop Options header would ignore the Hop-by-Hop options header and forward the packet only based on the destination IP address. Network nodes which support Hop-by-Hop Options header, but do not support the NR option would ignore the NR option and forward the packet only based on the destination IP address. The network node may process the rest of the Hop-by-Hop options in the Hop-by-Hop Options header.

4. Operational Considerations

As described in section 4.8 of [RFC8200], network nodes may be configured to ignore the Hop-by-Hop Options header, drop packets containing a Hop-by-Hop Options header, or assign packets containing a Hop-by-Hop Options header to a slow processing path. In networks with such network nodes, packets mapped to an NRP may be dropped due to the existence of the Hop-by-Hop Options header. Thus operators need to make sure that all the network nodes involved in an NRP can either process the Hop-by-Hop Options header in full forwarding rate, or ignore the Hop-by-Hop Options header. Since an NRP is associated with a logical network topology, one practical approach is to ensure that all the network nodes involved in that logical topology support the processing of the Hop-by-Hop Options header and the NR option in the fast path, and constrain the packet forwarding path to the logical topology of the NRP.

[RFC9673] specifies the modified procedures for the processing of IPv6 Hop-by-Hop Options header, with the purpose of making the Hop-by-Hop Options header useful. Network nodes complying with [RFC9673] will not drop packets with Hop-by-Hop Options header and the NR option.

5. Considerations about Generalization

This section gives some analysis about to what extent the semantics of NR Option could be generalized, and how the generalization could be achieved with the encoding specified in section 2.

Based on the NRP definition in [RFC9543], the concept of NRP could be extended as: an underlay network construct which is associated with a set of network-wide attributes and states maintained on each participating network node. The attributes associated with an NRP may include but not limited to, forwarding plane resources, network topologies, and network functions etc.

* The network resource can refer to various type of forwarding plane resources, including link bandwidth, buffering and queueing resources.

* The network resource can refer to topologies with multipoint-to-multipoint, point-to-point, point-to-multipoint, or multipoint-to-point connectivity.

* The network resources may include both packet forwarding actions and other types network functions which can be executed on data packets.

Thus the semantics of network resource can be quite generic. Although generalization is something good to have, it would be important to understand and identify the boundary of generalization. In this document, it is anticipated that for one network attribute to be considered as network resource, it needs to be network-wide attribute rather than a single node-specific attribute. Thus whether a network-wide view can be provided or not could be considered as one prerequisite of making one attribute part of the NR option.

The format of the NR option contains the Flags field, the Context Type field, and the Unassigned field, which provide the capability for future extensions. That said, since the NR option needs to be processed by network nodes with full forwarding rate, the capability of network devices need to be considered when new semantics and encoding are introduced.

6. IANA Considerations

This document requests IANA to assign a new option type from "Destination Options and Hop-by-Hop Options" registry [IANA-HBH].

Hex Value	Binary Value act chg rest	Description	Reference
TBA	00 0 tba	NR Option	[this document]

This document requests IANA to create a new registry for the "NR Option Context Type" under the "Internet Protocol Version 6 (IPv6) Parameters" registry. The allocation policy of this registry is "Standards Action". The initial code points are assigned by this document as follows:

Value	Description	Reference
0	NRP Selector ID	[this document]
1-254	Unassigned	
255	Reserved	[this document]

This document requests IANA to create a new registry for the "NR Option Flags" under the "Internet Protocol Version 6 (IPv6) Parameters" registry. The allocation policy of this registry is "Standards Action". The initial code points are assigned by this document as follows:

Bit	Description	Reference
0	Strict Match	[this document]
1-7	Unassigned	

7. Security Considerations

The security considerations with IPv6 Hop-by-Hop Options header are described in [RFC8200], [RFC7045], [RFC9098] [RFC9099] and [RFC9673]. This document introduces a new IPv6 Hop-by-Hop option which is either processed in the fast path or ignored by network nodes, thus it does not introduce additional security issues.

8. Contributors

Zhibo Hu
Email: huzhibo@huawei.com

Lei Bao
Email: baolei7@huawei.com

9. Acknowledgements

The authors would like to thank Juhua Xu, James Guichard, Joel Halpern, Tom Petch, Aijun Wang, Zhenqiang Li, Tom Herbert, Adrian Farrel, Eric Vyncke, Erik Kline, Mohamed Boucadair, Ketan Talaulikar and Vishnu Pavan Beeram for their review and valuable comments.

10. References

10.1. Normative References

[IANA-HBH] "IANA, "Destination Options and Hop-by-Hop Options"", 2016, <https://www.iana.org/assignments/ipv6-parameters/>.

[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>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[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>.

10.2. 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-07, 2 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-07>.

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual Private Network (VPN) Terminology", RFC 4026, DOI 10.17487/RFC4026, March 2005, <https://www.rfc-editor.org/info/rfc4026>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <https://www.rfc-editor.org/info/rfc7045>.

[RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. Liu, "Operational Implications of IPv6 Packets with Extension Headers", RFC 9098, DOI 10.17487/RFC9098, September 2021, <https://www.rfc-editor.org/info/rfc9098>.

[RFC9099] Vyncke, ., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/info/rfc9099>.

[RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <https://www.rfc-editor.org/info/rfc9350>.

[RFC9673] Hinden, R. and G. Fairhurst, "IPv6 Hop-by-Hop Options Processing Procedures", RFC 9673, DOI 10.17487/RFC9673, October 2024, <https://www.rfc-editor.org/info/rfc9673>.

[TS23501] "3GPP TS23.501", 2016, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

Authors' Addresses

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

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

Chongfeng Xie
China Telecom
China Telecom Beijing Information Science & Technology, Beiqijia
Beijing
102209
China
Email: xiechf@chinatelecom.cn

Chenhao Ma
China Telecom
China Telecom Beijing Information Science & Technology, Beiqijia
Beijing
102209
China
Email: machh@chinatelecom.cn

Gyan Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com