

IDR
Internet-Draft
Intended status: Standards Track
Expires: 8 March 2026

R. Chen
ZTE Corporation
J. Dong
Huawei
D. Zhao
ZTE Corporation
L. Gong
China mobile
Y. Zhu
China Telecom
R. Pang
China Unicom
4 September 2025

SR Policies Extensions for Network Resource Partition in BGP-LS
draft-ietf-idr-bgp-ls-sr-policy-nrp-02

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 ID is an important network resource attribute associated with the Segment Routing (SR) policy and needs to be reported to the external components. This document defines a new TLV which enables the headend to report the NRP which the SR Policy Candidate Path (CP) is associated with using Border Gateway Protocol Link-State (BGP-LS).

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 8 March 2026.

Copyright Notice

Copyright (c) 2025 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 TLV in BGP-LS	3
3. Scalability Considerations	4
4. Acknowledgements	5
5. IANA Considerations	5
6. Security Considerations	5
7. References	5
7.1. Normative References	5
7.2. Informative References	6
Authors' Addresses	7

1. Introduction

Segment Routing (SR) Policy [RFC9256] is an ordered list of segments (i.e. instructions) that represent a source-routed policy. Packet flows are steered into a SR Policy on a node where it is instantiated called a headend node. The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy.

[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.

[I-D.ietf-teas-nrp-scalability] introduces a scalable data plane approach to support Network Slice is to carry a dedicated NRP ID in the data packet to identify the NRP the packet belongs to, so that the packet can be processed and forwarded using the subset of network resources allocated to the NRP. For SR Policy with IPv6 data plane,

the approach to encapsulate the NRP ID in IPv6 Hop-by-Hop Options header is defined in [I-D.ietf-6man-enhanced-vpn-vtn-id]. For SR Policy with MPLS data plane, one approach to encapsulate the NRP ID to the packet is defined in [I-D.ietf-mppls-mna-nrp-selector].

[I-D.ietf-idr-sr-policy-nrp] defines the extensions to BGP SR policy to specify the NRP which the SR Policy candidate path is associated with.

[I-D.ietf-idr-bgp-ls-sr-policy] describes a mechanism to distribute SR policy information to external components using BGP-LS. SR policy information can be used by external components for path computation, re-optimization, service placement, network visualization, etc.

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. The SR policy indicates the forwarding path of the data packet, and the NRP ID indicates the reserved resources along the path specified by the SR policy. By associating the SR policy with a specific NRP, the forwarding path and resource reservation along the path are ensured. BGP-LS reports the association between SR Policy Candidate Path (CP) and NRP-ID, which is used to synchronize the association information between the SR Policy CP and the NRP-ID.

This document defines a new TLV which enables the headend to report the NRP which the SR Policy CP is associated with by using BGP-LS.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Carrying NRP TLV in BGP-LS

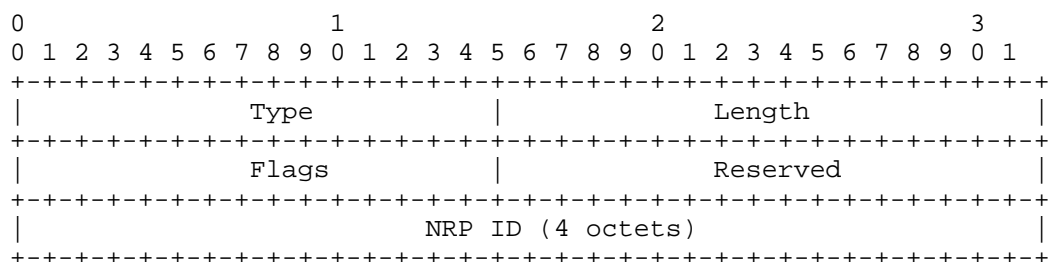
A SR Policy CP MAY be instantiated with a specific NRP on the headend node via a local configuration, PCEP, or BGP SR Policy signaling. Then the state and attributes of the NRP associated with the candidate path of SR policy can be distributed to the controller.

In order to collect configuration and states of the NRP SR policy, this document defines a new SR Policy state TLV called the NRP TLV which enables the headend to report the state at the SR Policy CP level.

This TLV is carried in the optional non-transitive BGP Attribute "LINK_STATE Attribute" defined in [RFC9552] associated with the SR Policy CP NLRI type.

This TLV is optional and only one these TLVs is advertised for a given SR Policy CP. If multiple TLVs are present, then the first one is considered valid and the rest are ignored as describe in [I-D.ietf-idr-bgp-ls-sr-policy].

The TLV has the following format:



where:

Type: TBD.

Length: 8. The total length of the value field not including Type and Length fields.

Flag: 2-octet flag field. None is defined at this stage. The flags SHOULD be set to zero on transmission and MUST be ignored on receipt.

RESERVED: 2-octet reserved bits. It SHOULD be set to zero on transmission and MUST be ignored on receipt.

NRP ID: 4-octet domain significant identifier of Network Resource Partition. Value 0 and 0xFFFFFFFF are reserved. NRP ID is planned by network operator.

3. Scalability Considerations

The mechanism specified in this document defines a mechanism for the headend of the SR policy to report configuration and states of an SR policy carrying the NRP information by using BGP-LS. Although the proposed mechanism allows an NRP MAY support many SR policies, in normal case the association between the SR Policy and the NRP is considered to be 1:1. As the number of NRP increases, the number of

SR Policies would also increase accordingly, and the status reported by the headend increases accordingly. However, the relationship between the SR Policy and the NRP is relatively stable and does not change status frequently. On the other hand this will only cause an increase in the status reporting information of the head node, the impacts to the BGP control plane are considered acceptable.

4. Acknowledgements

The authors would like to thank Susan Hares, Joel Halpern, Ketan Talaulikar, Vishnu Pavan Beeram, Adrian Farrel, and Changwang Lin for their review and discussion of this document.

5. IANA Considerations

IANA maintains a registry group called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a registry called "BGP-LS NLRI and Attribute TLVs". The following TLV codepoints are suggested (for early allocation by IANA):

Codepoint	Description	Reference
TBD	NRP	This document

6. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [RFC4271] for a discussion of BGP security. Security considerations for acquiring and distributing BGP-LS information are discussed in [RFC9552]. Security considerations for acquiring and distributing BGP-LS SR Policy information are discussed in [I-D.ietf-idr-sr-policy-nrp].

Additionally, it MAY be considered that reporting the NRP ID in association with the SR policy constitutes a risk to confidentiality of mission-critical or commercially sensitive information about the network. It is the responsibility of the network operator to ensure that only trusted nodes (that include both routers and controller applications) within the SR domain are configured to receive such information.

7. References

7.1. Normative References

[I-D.ietf-idr-bgp-ls-sr-policy]

Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.

[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>>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.

[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>>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

7.2. Informative References

[I-D.ietf-6man-enhanced-vpn-vtn-id]

Dong, J., Li, Z., Xie, C., Ma, C., and G. S. Mishra, "Carrying Network Resource (NR) related Information in IPv6 Extension Header", Work in Progress, Internet-Draft, draft-ietf-6man-enhanced-vpn-vtn-id-12, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-enhanced-vpn-vtn-id-12>>.

[I-D.ietf-idr-sr-policy-nrp]

Dong, J., Hu, Z., and R. Pang, "BGP SR Policy Extensions for Network Resource Partition", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-nrp-03, 3 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-nrp-03>>.

[I-D.ietf-mppls-mna-nrp-selector]

Li, T., Drake, J., Beeram, V. P., Saad, T., and I. Meilik,
"MPLS Network Actions for Network Resource Partition
Selector", Work in Progress, Internet-Draft, draft-ietf-
mppls-mna-nrp-selector-00, 13 May 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-mppls-
mna-nrp-selector-00](https://datatracker.ietf.org/doc/html/draft-ietf-mppls-mna-nrp-selector-00)>.

[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](https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-scalability-07)>.

[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>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Jie Dong
Huawei
Beijing
China
Email: jie.dong@huawei.com

Detao Zhao
ZTE Corporation
Nanjing
China
Email: zhao.detao@zte.com.cn

Liyan Gong
China mobile
Beijing
China
Email: gongliyan@chinamobile.com

Yongqing Zhu
China Telecom
Guangzhou
China
Email: zhuyq8@chinatelecom.cn

Ran Pang
China Unicom
Beijing
China
Email: pangran@chinaunicom.cn