

SPRING
Internet-Draft
Intended status: Standards Track
Expires: 1 November 2025

W. Cheng
W. Jiang
China Mobile
R. Chen
ZTE Corporation
C. Lin
New H3C Technologies
G. Zhang
China Mobile
30 April 2025

Resource Guarantee for SRv6 Policies
draft-cheng-spring-srv6-policy-resource-guarantee-05

Abstract

This document defines a new SRv6 Endpoint behavior which can be used to associate with a set of network resource partition (e.g. bandwidth, buffer and queue resources) Programming, called End.NRP. By using the End.NRP SID to build its segment list, the SRv6 policy has the capability to program network resources and achieve strict SLA guarantees.

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 1 November 2025.

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
1.2. Terminology	3
2. End.NRP Behavior	3
3. Use Cases for End.NRP behavior	4
4. Security Considerations	6
5. Implementation Status	6
5.1. ZTE Corp	7
5.2. New H3C Technologies	8
6. IANA Considerations	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	10
Authors' Addresses	10

1. Introduction

The concept of Network Resource Partition is introduced in [RFC9543]. A Network Resource Partition (NRP) is a set of network resources that are allocated from the underlay network to carry a specific set of network traffic and meet the required SLOs and SLEs.

Segment Routing (SR) [RFC8402] leverages the source routing paradigm. An ingress node steers a packet through an ordered list of instructions, called "segments". Each one of these instructions represents a function to be called at a specific location in the network. A function is locally defined on the node where it is executed and may range from simply moving forward in the segment list to any complex user-defined behavior.

SR Policy is an ordered list of segments (i.e. instructions) that represent a source-routed policy. The packets steered into an SR Policy have an ordered list of segments associated with that SR Policy.

Since the SRv6 Endpoint behavior defined in [RFC8986] are not associated with a set of network resource partition of the interface for slices/slice aggregate(e.g.End.X just forwards to an endpoint with cross-connect to a 'layer-3 adjacency' or L2 bundles). Therefore, SRv6 policies can't achieve strict SLA guarantees.

[I-D.ietf-spring-resource-aware-segments] extends the SR paradigm by associating SIDs with network resource attributes. On the basis of [I-D.ietf-spring-resource-aware-segments], this document defines a new SRv6 Endpoint behavior which can be used to associate with a set of network resource partition (e.g. bandwidth, buffer and queue resources) Programming, called End.NRP. By using the End.NRP SID to build its segment list , the SRv6 policy has the capability to program network resources and achieve strict SLA guarantees.

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.

1.2. Terminology

The following terminology is used in this document:

- * NRP: Network Resource Partition
- * PSP: Penultimate Segment Pop
- * SLA: Service Level Agreement
- * SR: Segment Routing
- * SRv6:The instantiation of SR on the IPv6 data plane
- * USD: Ultimate Segment Decapsulation
- * USP: Ultimate Segment Pop

2. End.NRP Behavior

This section defines a new SRv6 Endpoint behavior which can be used to associate with a set of network resource partition (e.g. bandwidth, buffer and queue resources) Programming, called End.NRP. The End.NRP is a variant of the End.X behavior defined in [RFC8986].

Any SID instance of End.NRP behavior is associated with two sets: J1 and J2.

J1: one or more L3 adjacencies or L2 bundles
J2: NRP of J1

When N receives a packet destined to S and S is a local End.NRP SID, the line S15 from the End processing defined in [RFC8986] is replaced by the following:

S15. Submit the packet to the IPv6 module for transmission to the new destination via a member of J1, using the NRP identified by J2

This End.NRP SID SHOULD support the Penultimate Segment Pop (PSP) of the SRH, Ultimate Segment Pop (USP) of the SRH, and Ultimate Segment Decapsulation (USD) flavors defined in [RFC8986] either individually or in combinations. The SRH processing of the End.NRP behavior with PSP, USP, and USD is the same as [RFC8986].

This End.NRP SIDs can be allocated either by a centralized network controller or by the network nodes, and the End.NRP behavior MAY be announced using IGP or BGP-LS. The detailed protocol extension will be described in a separate document.

3. Use Cases for End.NRP behavior

This section describes possible procedures for the End.NRP behavior.

A group of End.NRP SIDs SHOULD be allocated for the set of network resources associated with the SRv6 Policies, so that different End.NRP SIDs SHOULD be used to steer service traffic into different set of link resources (e.g. bandwidth, buffer and queue resources) in packet forwarding.

Below is the possible procedures:

- * The controller get the topology information, calculate the SR Policy path based on SLA.
- * The controller cooperates with the network nodes to complete resource reservation and the End.NRP SID allocation along the SR Policy path.
- * The controller use the End.NRP SID to build the SID list for the explicit path.

- * Then the controller inform the headend the resource guaranteed path by various means including: via BGP [I-D.ietf-idr-sr-policy-safi], configuration or PCEP [RFC8664] [I-D.ietf-pce-segment-routing-policy-cp].

SRv6 Policy1 from PE1 to PE2
segment list <End.NRP11, End.NRP12, End.NRP13>

SRv6 Policy2 from PE1 to PE2
segment list <End.NRP21, End.NRP22, End.NRP23>




Figure 1 shows an example for the End.NRP behavior.

As shown in Figure 1, there are two customers with different leased line requirements from PE1 to PE2:

- * leased line1 : 1G BandWidth with strict SLA guarantee.
- * leased line2 : 2G BandWidth with strict SLA guarantee.

Below is the possible procedures:

- * The controller get the topology information, calculate the SRv6 Policy 1 and SRv6 Policy 2 based on SLA.
- * The controller cooperates with the network nodes to complete resource reservation and the End.NRP SID allocation along the SRv6 Policy1 and SRv6 Policy 2. Taking the interface PE1-P1 of SRv6 node PE1 along the SRv6 Policy 1 as an example, two different NRPs(e.g. two dedicated queues) are partitioned from the network resources of the physical link PE1-P1 (GE1/0/0).

- The NRP(Queue1:1G BW)of link PE1- P1 is reserved and associated with End.NRP11.
- The NRP(Queue2:2G BW)of link PE1- P1 is reserved and associated with End.NRP21.
- * The controller uses the End.NRP11, End.NRP12, and End.NRP13 to build the SID list for the SRv6 Policy1 and use the End.NRP21, End.NRP22, and End.NRP23 to build the SID list for the SRv6 Policy 2.
- * Then the controller inform the headend the segment list of SRv6 Policy 1 and the segment list of SRv6 Policy 2 by various means including: via BGP [I-D.ietf-idr-sr-policy-safi], configuration or PCEP [RFC8664] [I-D.ietf-pce-segment-routing-policy-cp].

The traffic from customer1 and customer2 will be forwarded to PE2 through the NRPs previously reserved for each hop link on the path of SRv6 Policy1 and SRv6 Policy2 respectively, thus Customer 1 and Customer 2 are provided with end-to-end 1G bandwidth resources and 2G bandwidth resources respectively, and private line services are guaranteed by strict SLAs.

4. Security Considerations

The security requirements and mechanisms described in [RFC8402], [RFC8754] and [RFC8986] also apply to this document.

This document introduces a new SRv6 Endpoint behavior for implementation on the nodes support network resource partition in the network. As such, this document does not introduce any new security considerations.

5. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC7942]].

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC7942], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

5.1. ZTE Corp

- * Organization: ZTE Corporation.
- * Implementation: ZTE's M6000 Series Routers implementation of END.NRP [ZTE-IMP].
- * Description: The feature of SRv6 END.NRP has been implemented in ZTE M6000 series routers and follows the definition and mechanism as defined in section 2 and Section 3 including all the "MUST" and "SHOULD" clauses.
- * Maturity Level: Beta
- * Coverage: ALL, section 2 and use case section 3.
- * Version: Draft-02
- * Licensing: N/A
- * Implementation experience: Nothing specific.
- * Contact: zhu.xiaolong@zte.com.cn
- * Last updated: October 21, 2023

5.2. New H3C Technologies

- * Organization: New H3C Technologies.
- * Implementation: H3C CR16000, CR19000 series routers implementation.
- * Description: Section 2 and Section 3 including all the "MUST" and "SHOULD" clauses have been implemented in above mentioned New H3C Products for testing.
- * Maturity Level: Beta
- * Coverage: ALL, section 2 and use case section 3.
- * Version: Draft-02
- * Licensing: N/A
- * Implementation experience: Nothing specific.
- * Contact: linchangwang.04414@h3c.com
- * Last updated: October 21, 2023

6. IANA Considerations

The document defines a new SRv6 Endpoint behavior called End.NRP.

This I-D requests the IANA to allocate, within the "SRv6 Endpoint Behaviors" sub-registry belonging to the top-level "Segment-routing with IPv6 dataplane (SRv6) Parameters" registry, the following allocations:

Value	Endpoint Behavior	Reference
TBD1	End.NRP	[This.ID]

7. Acknowledgements

The authors would like to thank Detao Zhao and Jie Dong for their review and comments.

8. References

8.1. Normative References

- [I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.
- [I-D.ietf-pce-segment-routing-policy-cp]
Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-27, 4 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-27>>.
- [I-D.ietf-spring-resource-aware-segments]
Dong, J., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-11, 2 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-resource-aware-segments-11>>.
- [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>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.
- [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>>.
- [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>>.

- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

8.2. Informative References

- [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>>.
- [ZTE-IMP] "ZTE M6000-S Routers", 21 October 2023, <https://www.zte.com.cn/china/product_index/ip_network/item02/zxr10-m6000-s/zxr10_m6000_s.html>.

Authors' Addresses

Weiqiang Cheng
China Mobile
Beijing
China
Email: chengweiqiang@chinamobile.com

Wenying Jiang
China Mobile
Beijing
China
Email: jiangwenying@chinamobile.com

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

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

Geng Zhang
China Mobile
China
Email: zhanggeng@chinamobile.com