

Workgroup: Network Working Group
Internet-Draft:
draft-zhang-sr-policy-enhanced-detnet-00
Published: 11 July 2022
Intended Status: Standards Track
Expires: 12 January 2023
Authors: L. Zhang X. Geng Z. Li
 Huawei Huawei Huawei
 SR Policy for enhanced DetNet

Abstract

SR Policy is a set of candidate SR paths consisting of one or more segment lists and necessary path attributes. It enables instantiation of an ordered list of segments with a specific intent for traffic steering. DetNet provides the capability to carry specified unicast or multicast data flows with extremely low data loss rates and bounded end-to-end latency within a network domain. This document defines the SR policy enhancement to carry the Bounded Latency Information with a candidate path of SR policy. So that BLI behavior can be enabled automatically when the SR Policy is applied.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 12 January 2023.

Copyright Notice

Copyright (c) 2022 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. Terminology and Conventions](#)
 - [2.1. Requirement Language](#)
 - [2.2. Terminology](#)
- [3. BLI Encoding in SR Policy](#)
 - [3.1. BLI List Sub-TLV](#)
 - [3.2. Shared BLI sub-TLV](#)
- [4. Procedures](#)
- [5. IANA Considerations](#)
- [6. Security Considerations](#)
- [7. Acknowledgements](#)
- [8. Normative References](#)
- [Authors' Addresses](#)

1. Introduction

Segment Routing Policy is defined in [[I-D.ietf-spring-segment-routing-policy](#)]. A SR Policy is a set of candidate path which consist of one or more segment lists. The headend node instructs the source routing and writes it into package. The packets steered into an SR Policy have an ordered list of segments associated with that SR Policy written into them. [[RFC8655](#)] provides the overall architecture for Deterministic Networking (DetNet), which provides the capability to carry specified unicast or multicast data flows with extremely low data loss rates and bounded end-to-end latency within a network domain. Based on this, [[I-D.ietf-detnet-bounded-latency](#)] proposed a timing model for sources, destinations, and DetNet transit nodes. Using the model, it provides a methodology to compute end-to-end latency and backlog bounds for various queuing methods. [[I-D.yzz-detnet-enhanced-data-plane](#)] enhances the DetNet data plane by introducing Bounded Latency Information (BLI) which facilitates DetNet transit nodes to guarantee the bounded latency transmission in data plane. Based on that, [[I-D.geng-spring-sr-enhanced-detnet](#)] defines how to leverage Segment Routing (SR) and Segment Routing over IPv6 (SRv6) to implement bounded latency. For An automatic network, the SR Policy with Bounded Latency Information

can facilitate the bounded latency transmission and enable the automation of SR service.

This document defines the SR policy enhancement to carry the Bounded Latency Information with a candidate path of SR policy. So that BLI behavior can be enabled automatically when the SR Policy is applied.

2. Terminology and Conventions

2.1. Requirement 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](#)].

2.2. Terminology

The abbreviations used in this document are:

BLI: Bounded Latency Information

SR: Segment Routing

SID: Segment Identifier


3. BLI Encoding in SR Policy

The BLI is proposed by [[I-D.yzz-detnet-enhanced-data-plane](#)] to facilitate DetNet transit nodes to guarantee the bounded latency transmission in data plane. In order to specify the bounded latency features that the candidate path is associated with, this document defines two types of new sub-TLV in the BGP Tunnel Encapsulation Attribute for SR Policy [[I-D.ietf-spring-segment-routing-policy](#)] for different scenarios.

3.1. BLI List Sub-TLV

When all of the nodes/adjacencies in the explicit path indicated by the segment list request different BLI to guarantee bounded latency, a BLI list sub-TLV is defined.

The BLI list sub-TLV is formatted as follows.



Where:

Type: to be assigned by IANA.

Length: 16 bits length value to indicate the length of BLI list in octet.

BLI List [1... m]: 64 bits length BLI structure, representing the nth BLI in the BLI list.

The BLI in the BLI List corresponds to the Segment in the Segment List one by one. The length of the BLI List depends on the num of Segment in the Segment List.

The encoding structure of BGP SR Policy with the BLI list sub-TLV is expressed as below:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encaps Attribute (23)

Tunnel Type: SR Policy

Binding SID

Preference

Priority

Policy Name

Explicit NULL Label Policy (ENLP)

Segment List

BLI List

Weight

Segment

Segment


...

...

3.2. Shared BLI sub-TLV

When all of the nodes/adjacencies in the explicit path indicated by the segment list request BLI to guarantee bounded latency with the same BLI value, the Shared BLI sub-TLV is defined.

The Shared BLI sub-TLV is defined as follows:



Where:

Type: to be assigned by IANA.

Length: 16 bits value indicate the length of BLI.

BLI: 64 bits value of Bounded Latency Information to guarantee the bounded latency, the format of it is defined in section 3.1.

The encoding structure of BGP SR Policy with the Per-segment BLI sub-TLV is expressed as below:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encaps Attribute (23)

Tunnel Type: SR Policy

Binding SID

Preference

Priority

Policy Name

Explicit NULL Label Policy (ENLP)

Segment List

Shared BLI

Weight

Segment

Segment

• • •

■ ■ ■

4. Procedures

When a candidate path of SR Policy is a bounded-latency routing path, the originating node of SR policy MUST include the associated

bounded latency information in the BGP Tunnel Encapsulation Attribute of the BGP SR Policy. The other fields and attributes in BGP SR Policy should follow the mechanism as defined in [[I-D.ietf-idr-segment-routing-te-policy](#)].

When a BGP speaker receives an SR Policy which is acceptable and usable according to the rules as defined in [[I-D.ietf-idr-segment-routing-te-policy](#)], and the SR Policy candidate path selected as the best candidate path is a bounded-latency path, the receiver node of the SR Policy MUST encapsulate the specific bounded latency information to the header of packets steered to the SR Policy. For SR Policy with IPv6 data plane and MPLS data plane, the possible approach is to encapsulate the BLI to the packet using the mechanism defined in [[I-D.yzz-detnet-enhanced-data-plane](#)] and [[I-D.geng-spring-sr-enhanced-detnet](#)].

5. IANA Considerations

IANA is requested to make the assignment from the "BGP Tunnel Encapsulation Attribute sub-TLVs" registry as follows.

Value	Name	Reference
TBD1	BLI List sub-TLV	This document
TBD2	Shared BLI sub-TLV	This document

6. Security Considerations

TBD

7. Acknowledgements

8. Normative References

[[I-D.geng-spring-sr-enhanced-detnet](#)] Geng, X., Li, Z., and T. Zhou, "Segment Routing for Enhanced DetNet", Work in Progress, Internet-Draft, draft-geng-spring-sr-enhanced-detnet-00, 11 July 2022, <<https://www.ietf.org/archive/id/draft-geng-spring-sr-enhanced-detnet-00.txt>>.

[[I-D.ietf-detnet-bounded-latency](#)]

Finn, N., Boudec, J. L., Mohammadpour, E., Zhang, J., and B. Varga, "DetNet Bounded Latency", Work in Progress, Internet-Draft, draft-ietf-detnet-bounded-latency-10, 8 April 2022, <<https://www.ietf.org/archive/id/draft-ietf-detnet-bounded-latency-10.txt>>.

[I-D.ietf-idr-segment-routing-te-policy]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-18, 16 June 2022, <<https://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-18.txt>>.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-22, 22 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-22.txt>>.

[I-D.yzz-detnet-enhanced-data-plane] Yang, F., Zhou, T., and L.

Zhang, "DetNet Enhanced Data Plane", Work in Progress, Internet-Draft, draft-yzz-detnet-enhanced-data-plane-00, 10 July 2022, <<https://www.ietf.org/archive/id/draft-yzz-detnet-enhanced-data-plane-00.txt>>.

[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>>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <<https://www.rfc-editor.org/info/rfc8655>>.

Authors' Addresses

Li Zhang
Huawei

Email: zhangli344@huawei.com

Xuesong Geng
Huawei

Email: gengxuesong@huawei.com

Zhenbin Li
Huawei

Email: lizhenbin@huawei.com