

Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 3, 2019

C. Li
Z. Li
Huawei Technologies
July 02, 2018

Segment Routing Path MTU in BGP
draft-li-idr-sr-policy-path-mtu-00

Abstract

Segment routing is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR policy is a set of candidate SR paths consisting of one or more segment lists with necessary path attributes. However, the path maximum transmission unit (MTU) information for SR path is not available in the SR policy since the SR does not require signaling. This document defines extensions to BGP to distribute path MTU information within SR policies.

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 January 3, 2019.

Copyright Notice

Copyright (c) 2018 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	SR Policy for Path MTU	3
3.1.	SR Path MTU Sub-TLV	4
4.	Operations	4
5.	IANA Considerations	5
6.	Security Considerations	5
7.	Acknowledgements	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	6
	Authors' Addresses	6

[1.](#) Introduction

Segment routing (SR) [[I-D.ietf-spring-segment-routing](#)] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. The ingress node steers packets into a specific path according to the Segment Routing Policy (SR Policy) as defined in [[I-D.ietf-spring-segment-routing-policy](#)]. For distributing SR policies to the headend, [[I-D.ietf-idr-segment-routing-te-policy](#)] specifies a mechanism by using BGP, and new sub-TLVs are defined for SR Policies in BGP UPDATE message.

The maximum transmission unit (MTU) is the largest size packet or frame, in bytes, that can be sent in a network. An MTU that is too large might cause retransmissions. Too small an MTU might cause the router to send and handle relatively more header overhead and acknowledgments. When an LSP is created across a set of links with different MTU sizes, the ingress router needs to know what the smallest MTU is on the LSP path. If this MTU is larger than the MTU of one of the intermediate links, traffic might be dropped, because MPLS packets cannot be fragmented. Also, the ingress router may not be aware of this type of traffic loss, because the control plane for

the LSP would still function normally. [[RFC3209](#)] specify the mechanism of MTU signaling in RSVP.

However, the path maximum transmission unit (MTU) information for SR path is not available since the SR does not require signaling. This document defines extensions to BGP to distribute path MTU information within SR policies. The MTU information can be obtained via IGP [[I-D.zhu-idr-bgp-ls-path-mtu](#)], BGP-LS [[I-D.hu-lsr-isis-path-mtu](#)] or some other means.

2. Terminology

This memo makes use of the terms defined in [[I-D.ietf-spring-segment-routing](#)] and [[RFC3209](#)].

3. SR Policy for Path MTU

As defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] , the SR Policy Encoding structure is as follows:

```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
  Attributes: Tunnel Encaps Attribute (23)
  Tunnel Type: SR Policy
    Binding SID
    Preference
    Segment List
      Weight
      Segment
      Segment
      ...
```

As introduced in [Section 1](#), each SR path may have a path MTU, an SR policy carrying a SR path MTU is expressed as below:

```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
  Attributes: Tunnel Encaps Attribute (23)
  Tunnel Type: SR Policy
    Binding SID
    Preference
    Segment List
      Weight
      Path MTU
      Segment
      Segment
      ...
```


3.1. SR Path MTU Sub-TLV

This section defines an SR Path MTU sub-TLV, and it is included in the segment list sub-TLV.

An SR Path MTU sub-TLV is associated with an SR path specified by a segment list sub-TLV or path ID as defined in [\[I-D.cheng-spring-mpls-path-segment\]](#) and [\[I-D.li-spring-passive-pm-for-srv6-np\]](#), and it MUST appear only once within a Segment List sub-TLV. It has the following format:




Figure 1. Path MTU sub-TLV

Where:

Type: to be assigned by IANA (suggested value 11).

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

Reserved: 16 bits reserved and MUST be set to 0 on transmission and MUST be ignored on receipt.

Path MTU: 4 bytes value of Path MTU. The value can be calculated by a central controller or other devices based on the information that learned via IGP of BGP-LS or other means.

Whenever the path MTU of a physical or logical interface is changed, a new SR policy with new path MTU information should be updated accordingly by BGP.

4. Operations

The document does not bring new operation beyond the description of operations defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#). The existing operations defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#) can apply to this document directly.

Typically but not limit to, the SR policies carrying path MTU information are configured by a controller.

After configuration, the SR policies carrying path MTU information will be advertised by BGP update messages. The operation of advertisement is the same as defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#), as well as the reception.

The consumer of the SR policies is not the BGP process. The operation of sending information to consumers is out of scope of this document.

5. IANA Considerations

TBA

6. Security Considerations

TBA

7. Acknowledgements

TBA

8. References

8.1. Normative References

- [I-D.cheng-spring-mpls-path-segment]
Cheng, W., Wang, L., Li, H., Chen, M., Zigler, R., and S. Zhan, "Path Segment in MPLS Based Segment Routing Network", [draft-cheng-spring-mpls-path-segment-01](#) (work in progress), March 2018.
- [I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Jain, D., Mattes, P., Rosen, E., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-03](#) (work in progress), May 2018.
- [I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-15](#) (work in progress), January 2018.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., daniel.voyer@bell.ca, d., bogdanov@google.com, b., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-01](#) (work in progress), June 2018.

[I-D.li-spring-passive-pm-for-srv6-np]

Li, C. and M. Chen, "Passive Performance Measurement for SRv6 Network Programming", [draft-li-spring-passive-pm-for-srv6-np-00](#) (work in progress), March 2018.

[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>>.

8.2. Informative References

[I-D.hu-lsr-isis-path-mtu]

Hu, Z., Zhu, Y., Li, Z., and L. Dai, "IS-IS Extensions for Path MTU", [draft-hu-lsr-isis-path-mtu-00](#) (work in progress), June 2018.

[I-D.zhu-idr-bgp-ls-path-mtu]

Zhu, Y., Hu, Z., Yan, G., and J. Yao, "BGP-LS Extensions for Advertising Path MTU", [draft-zhu-idr-bgp-ls-path-mtu-00](#) (work in progress), June 2018.

[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>>.

Authors' Addresses

Cheng Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: chengli13@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

