

Workgroup: Network Working Group
Internet-Draft: draft-chen-idr-mbinding-03
Published: 10 November 2023
Intended Status: Standards Track
Expires: 13 May 2024
Authors: H. Chen B. Decraene G. Mishra Y. Fan
 Futurewei Orange Verizon Casa Systems
 A. Wang X. Liu
 China Telecom Alef Edge
 BGP for Mirror Binding

Abstract

BGP is used to distribute a binding to a node. The binding includes a binding SID and a path represented by a list of SIDs. This document describes extensions to BGP for distributing the information about the binding to a protecting node. For an SR path via the node with the binding SID, when the node fails, the protecting node such as the upstream neighbor on the path uses the information to protect the binding SID of the failed node.

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.

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 13 May 2024.

Copyright Notice

Copyright (c) 2023 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. Extensions to BGP](#)
- [3. Procedure for Updating Information](#)
- [4. Security Considerations](#)
- [5. IANA Considerations](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

[[I-D.ietf-idr-segment-routing-te-policy](#)] specifies how BGP may be used to distribute a Segment Routing (SR) Policy to a node in a network. An SR Policy may contain a binding, which includes a binding SID and a path represented by a list of SIDs.

After a BGP as a controller distributes the binding to the node, the node forwards the packet with the binding SID according to the first SID in the list. It replaces the binding SID in the packet with the list of SIDs and forwards the packet using the FIB entry for the top SID (i.e., the first SID) in the packet.

When the node fails, suppose that a protecting node (e.g., the upstream neighbor as PLR) of the node has the corresponding binding protection information for protecting the binding SID of the node. The information includes the binding SID, the list of SIDs and an identifier of the node. After the upstream neighbor as PLR detects the failure of the node, for a packet with the node SID of the failed node received, it protects the binding SID of the failed node. It replaces the binding SID in the packet with the list of SIDs, forwards the packet without going through the failed node

towards the top SID (i.e., the first SID, assuming it is a node SID for simplicity here). The requirement for distributing the binding protection information to a protecting node and how a protecting node protect the binding SID of the failed node are out of scope of this document and described in [\[I-D.ietf-spring-segment-protection-sr-te-paths\]](#) and [\[I-D.hu-spring-segment-routing-proxy-forwarding\]](#). (Note: the second reference will be removed after the first one includes enough text for protecting binding SIDs of a node).

This document specifies some procedures and extensions to BGP for distributing the information to a protecting node that may protect the failed node.

2. Extensions to BGP

This section defines a new Binding Protection sub-TLV under a Tunnel Encapsulation Attribute TLV of type 15 (i.e., SR Policy TLV). A Tunnel Encapsulation Attribute contains a Tunnel Encapsulation Attribute TLV. The structure containing a Binding Protection sub-TLV is shown below.

Attributes:

- Tunnel Encapsulation Attribute (23)
 - Tunnel Type (15): SR Policy TLV
 - Preference sub-TLV
 - Binding SID sub-TLV
 - SRv6 Binding SID sub-TLV
 - Explicit NULL Label Policy (ENLP) sub-TLV
 - Priority sub-TLV
 - Policy Candidate Path Name sub-TLV
 - Policy Name sub-TLV
 - Binding Protection sub-TLV
 - Segment List sub-TLV
 - Weight sub-TLV
 - Segment sub-TLV
 - Segment sub-TLV
 - ...
 - ...

The format of a Binding Protection sub-TLV is illustrated below.




Figure 1: Binding Protection sub-TLV Format

Type: Its value (TBD1) is to be assigned by IANA.

Length: It is variable.

Flags: 1 octet of flags. No flags is defined now. MUST be set to zero by the sender and MUST be ignored by the receiver.

sub-TLVs: This field contains the sub-TLV below to indicate the node to be protected (i.e., the protected node).

- o Protected Node ID sub-TLV indicating the ID of the Protected Node.

When an SR Policy (i.e., SR Policy TLV) contains a binding SID and a path with a protected node, the SR policy is for distributing the binding protection information. The binding SID is encoded by a Binding SID sub-TLV or SRv6 Binding SID sub-TLV, the path is encoded by a Segment List Sub-TLV, and the node is encoded by a Binding Protection sub-TLV.

When an SR Policy contains a binding SID and a path without a protected node, the SR policy is for replacing the Binding SID with the path (i.e., the list of SIDs) when the node receives a packet with the Binding SID.

The format of Protected Node ID sub-TLV is illustrated below.




Figure 2: Protected Node ID sub-TLV Format

Type:

Its value (1) indicates the type of Protected Node ID sub-TLV.

Length: Its value (4) indicates the length of the value field of the sub-TLV is 4.

Protected Node ID: 4-octet field contains the identifier (ID) of the Protected Node.

3. Procedure for Updating Information

When a BGP sends a binding to node N for a SR path going through node N, the BGP distributes the corresponding binding protection information to the possible protecting nodes. The upstream neighbor of node N along the SR path is a protecting node. If node N is a loose hop in the SR path with N's node ID, the upstream node having its node SID next to N's node ID in the SR path and the neighbors of node N are the possible protecting nodes.

After sending the binding to node N and distributing the corresponding binding protection information to the protecting nodes, if BGP removes the binding from node N, it removes the corresponding binding protection information from the protecting nodes. If BGP changes the binding in node N, it changes the corresponding binding protection information in the protecting nodes.

4. Security Considerations

Protocol extensions defined in this document do not affect the BGP security other than those as discussed in the Security Considerations section of [\[RFC9012\]](#).

5. IANA Considerations

This document requests assigning a new sub-TLV in the registry "BGP Tunnel Encapsulation Attribute sub-TLVs" as follows:

Code Point	Description	Reference
TBD1	Binding Protection	[This document]

6. References

6.1. Normative References

[I-D.ietf-spring-segment-protection-sr-te-paths]

Hegde, S., Bowers, C., Litkowski, S., Xu, X., and F. Xu, "Segment Protection for SR-TE Paths", Work in Progress,

Internet-Draft, draft-ietf-spring-segment-protection-sr-te-paths-05, 27 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-protection-sr-te-paths-05>>.

- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

6.2. Informative References

[I-D.hu-spring-segment-routing-proxy-forwarding]

Hu, Z., Chen, H., Yao, J., Bowers, C., Zhu, Y., and Y. Liu, "SR-TE Path Midpoint Restoration", Work in Progress, Internet-Draft, draft-hu-spring-segment-routing-proxy-forwarding-24, 21 August 2023, <<https://datatracker.ietf.org/doc/html/draft-hu-spring-segment-routing-proxy-forwarding-24>>.

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

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-segment-routing-te-policy-26, 23 October 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-26>>.

Acknowledgments

The authors would like to thank Andrew Stone, Robert Raszuk, Ketan Talaulikar, Donald Eastlake, Zhibo Hu, Haibo Wang, Yao Liu, Changwang Lin, Jie Dong and Zhenqiang Li for their comments to this work.

Authors' Addresses

Huaimo Chen
Futurewei
Boston, MA,
United States of America

Email: huaimo.chen@futurewei.com

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Gyan S. Mishra
Verizon
13101 Columbia Pike
Silver Spring, MD 20904
United States of America

Phone: [301 502-1347](tel:3015021347)
Email: gyan.s.mishra@verizon.com

Yanhe Fan
Casa Systems
United States of America

Email: yfan@casa-systems.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China

Email: wangaj3@chinatelecom.cn

Xufeng Liu
Alef Edge
United States of America

Email: xufeng.liu.ietf@gmail.com