Workgroup: Network Working Group
Internet-Draft: draft-chen-idr-mbinding-00
Published: 23 October 2022
Intended Status: Standards Track
Expires: 26 April 2023
Authors: H. Chen (Futurewei)
B. Decraene (Orange)
G. Mishra (Verizon)
Y. Fan (Casa Systems)
A. Wang (China Telecom)
X. Liu (IBM Corporation)

BGP for Mirror Binding

Abstract

BGP is used to distribute a binding SID with a list of SIDs to a node. This document describes extensions to BGP for distributing the binding SID with the list of SIDs and an identifier of the node. When detecting the failure of the node, an neighbor of the node protects the binding SID of the failed node.

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 26 April 2023.

Copyright Notice

Table of Contents

1. Introduction
2. Extensions to BGP
3. Procedure for Updating Information
4. Security Considerations
5. IANA Considerations
   5.1. Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs
6. References
   6.1. Normative References
   6.2. Informative References
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 network node. An SR Policy is an ordered list of segments (i.e., instructions) that represent a source-routed policy. An SR Policy consists of one or more candidate paths, each consisting of one or more segment lists. An SR Policy may contain a binding SID associated with a path represented by a segment list (i.e., a list of SIDs).

After a BGP as a controller distributes an SR policy containing a binding SID associated with a list of SIDs to a network node, each neighbor of the node needs the information about the binding SID for protecting the binding SID of the node when the node fails. The information includes the binding SID, the list of SIDs and the identifier (ID) of the node. This document proposes some procedures and extensions to BGP for distributing the information.

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
                 ...
             ...
```

A Binding Protection sub-TLV contains the information about a binding SID of a network node for a neighbor of the node to protect the Binding SID of the node when the neighbor detects the failure of the node.

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

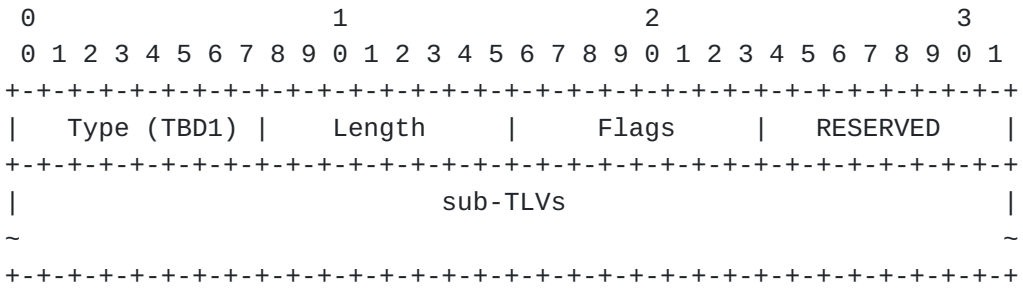

Figure 1: Binding Protection sub-TLV Format

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

When a 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 information of the node for protecting the binding SID of the node when the node fails (binding protection for short). 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 a 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 (binding for short).

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

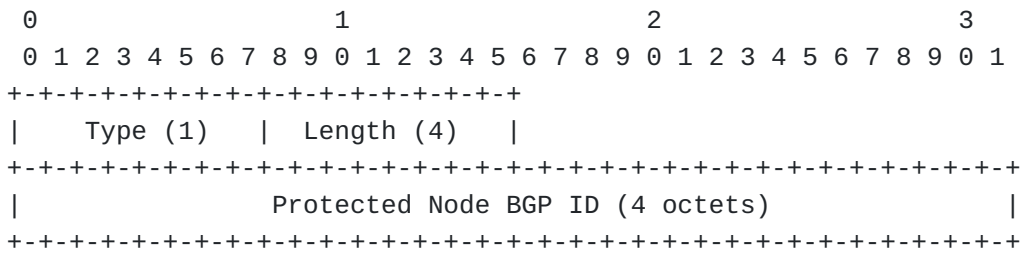

Figure 2: Protected Node BGP ID sub-TLV Format

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

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

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

3. Procedure for Updating Information

When a BGP sends a piece of binding information to node N in a first Update message, the BGP sends the corresponding binding protection information to each neighbor of node N in a second Update message. The first message contains a first SR Policy carried in MP_REACH_NLRI. The first SR Policy includes a binding SID and a path but does not include node N as a protected node. The second message contains a second SR Policy carried in MP_REACH_NLRI. The second SR Policy includes the binding SID, the path and node N as a protected node.

After a BGP sends the binding information to node N, if BGP removes the binding information from node N through sending a third Update message to node N, BGP removes the corresponding binding protection

information from each neighbor of node N through sending a fourth Update message to the neighbor. The third message contains a third SR Policy carried in MP_UNREACH_NLRI. The third SR Policy includes the binding SID and the path but does not include node N as a protected node. The fourth message contains a fourth SR Policy carried in MP_UNREACH_NLRI. The fourth SR Policy includes the binding SID, the path and node N as a protected node.

After a BGP sends the binding information to node N, if the BGP changes the binding information in node N through sending a fifth Update message to node N, BGP changes the corresponding binding protection information in each neighbor of node N through sending a sixth Update message to the neighbor. The fifth message contains a fifth SR Policy carried in MP_REACH_NLRI. The fifth SR Policy includes the binding SID and a (changed) path but does not include node N as a protected node. The sixth message contains a sixth SR Policy carried in MP_REACH_NLRI. The sixth SR Policy includes the binding SID, the (changed) path and node N as a protected node.

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 [RFC5575].

5. IANA Considerations

5.1. Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs

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

[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>.

[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.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-20, 27 July 2022, <https://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-20.txt>.

[RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", RFC 5575, DOI 10.17487/RFC5575, August 2009, <https://www.rfc-editor.org/info/rfc5575>.

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
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
IBM Corporation
United States of America

Email: xufeng.liu.ietf@gmail.com