

Inter-Domain Routing
Internet Draft
Intended status: Standard Track
Expires: September 27, 2025

Yisong Liu
China Mobile
C. Lin
New H3C Technologies
H. Bidgoli
Nokia
Z. Zhang
ZTE Corporation
March 27, 2025

Distribution of SR P2MP Policies and State using BGP-LS
draft-liu-idr-bgpls-sr-p2mp-policy-distribution-05

Abstract

This document specifies the extensions to BGP Link State (BGP-LS) to distribute SR P2MP Policies and state. This allows operators to establish a consistent view of the underlying multicast network state, providing an efficient mechanism for the advertisement and synchronization of SR P2MP policies.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on September 27, 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 (<http://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	3
1.1. Requirements Language	3
2. P2MP Policy NLRI Type	3
3. P2MP Policy Descriptor	6
3.1. P2MP Policy Candidate Path Descriptor	6
3.2. P2MP Policy Instance Descriptor	7
3.3. P2MP Policy downstream node Descriptor	8
4. P2MP Policy States TLV	10
4.1. SR P2MP Policy Name TLV	10
4.2. SR P2MP Policy Candidate Path Name TLV	10
4.3. SR P2MP Policy Candidate Path State TLV	11
4.4. Leaf List TLV	12
4.4.1. Leaf Sub-TLV	13
4.5. Path Instance TLV	14
4.5.1. Instance id Sub-TLV	15
4.6. Replication Segment TLV	16
4.7. Downstream List TLV	17
4.7.1. Downstream node Sub-TLV	19
4.8. Downstream-node Segment list TLV	20
4.8.1. Segment Sub-TLV	22
5. IANA Considerations	22
6. Security Considerations	23
7. References	23
7.1. Normative References	23
Contributors	24
Authors' Addresses	25

1. Introduction

As described in [RFC7752], BGP Link State (BGP-LS) provides a mechanism by which link-state and TE information can be collected from networks and shared with external components using the BGP routing protocol.

[draft-ietf-pim-sr-p2mp-policy] describes architecture to construct a Point-to-Multipoint (P2MP) tree to deliver Multi-point services in a Segment Routing domain.

[draft-ietf-idr-sr-p2mp-policy] defines the extension of BGP, which is used for the controller to advertise P2MP Policy to a group of nodes.

SR P2MP policies are set of policies that enable architecture for P2MP service delivery. A P2MP policy consists of candidate paths that connect the Root of the Tree to a set of Leaves. The P2MP policy is composed of replication segments.

This document defines the extension of BGP-LS, which is used to collect the state information of SR P2MP Policy on nodes. This information can be used by external components for path computation, re-optimization, service placement, network visualization, etc.

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.

2. P2MP Policy NLRI Type

The "Link-State NLRI" defined in [RFC7752] is extended to carry the P2MP Policy information. BGP speakers that wish to exchange P2MP Policy information MUST use the BGP Multiprotocol Extensions Capability Code(1) to advertise the corresponding (AFI, SAFI) pair, as specified in [RFC4760]. New TLVs carried in the Link_State Attribute defined in [RFC7752] are also defined to carry the attributes of a P2MP Policy in the subsequent sections.

The format of "Link-State NLRI" is defined in [RFC7752] as follows:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               NLRI Type                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Total NLRI Length                       |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Link-State NLRI (variable)              |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

A new NLRI Type needs to be defined to identify P2MP Policy information.

- o SR P2MP Policy Candidate Path NLRI (value TBD1)

The general format of the SR P2MP Policy NLRI defined in this document is as follows:

```

      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Protocol-ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Identifier                               |
|                               (64 bits)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Node Descriptor TLV                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               P2MP Policy Descriptors (variable)      //
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Where

- o Protocol-ID: Protocol-ID field specifies the component that owns the P2MP Policy state in the advertising node

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| Protocol-ID | NLRI information source protocol |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      TBD2   |      P2MP Policy                 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- o "Identifier" is an 8 octet value as defined in [RFC7752].
- o "Local Node Descriptor" (TLV 256) as defined in [RFC7752] that describes the advertising node.

- o "P2MP Policy Descriptors" consists of one or more of the TLVs listed as below for use with the P2MP Policy NLRI type advertisement.

Codepoint	Descriptor TLVs
TBD3	P2MP Policy Candidate Path
TBD4	P2MP Policy Instance
TBD5	P2MP Policy downstream node

The Local Node Descriptor TLV MUST include the following Node Descriptor TLVs:

- o BGP Router-ID (TLV 516) [RFC9086], which contains a valid BGP Identifier of the local node.
- o Autonomous System Number (TLV 512) [RFC7752], which contains the ASN or AS Confederation Identifier (ASN) [RFC5065], if confederations are used, of the local node.

The Local Node Descriptor TLV SHOULD include at least one of the following Node Descriptor TLVs:

- o IPv4 Router-ID of Local Node (TLV 1028) [RFC7752], which contains the IPv4 TE Router-ID of the local node when one is provisioned.
- o IPv6 Router-ID of Local Node (TLV 1029) [RFC7752], which contains the IPv6 TE Router-ID of the local node when one is provisioned.

The Local Node Descriptor TLV MAY include the following Node Descriptor TLVs:

- o BGP Confederation Member (TLV 517) [RFC9086], which contains the ASN of the confederation member (i.e. Member-AS Number), if BGP confederations are used, of the local node.
- o Node Descriptors as defined in [RFC7752].

3. P2MP Policy Descriptor

This section defines the P2MP Policy Descriptors TLVs which are used to describe the P2MP Policy being advertised by using the NLRI types defined in Section 2.

3.1. P2MP Policy Candidate Path Descriptor

The P2MP Policy Candidate Path Descriptor TLV identifies a Segment Routing P2MP Policy Candidate Path (CP) as defined in [draft-ietf-idr-sr-p2mp-policy]. This descriptor is used to report the state information of the P2MP Policy candidate path.

The TLV has the following format:

```

0                               1                               2                               3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|          Type                      |          Length                    |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Flags                     |          Reserved                   |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                                Root-ID (4 or 16 octets)                                //
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                Tree-ID (4 octets)                                |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                Distinguisher (4 octets)                        |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

- o Type: TBD6
- o Length: variable (valid values are 16, or 28)
- o Flags: 2-octet field with following bit positions defined. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

```

0                               1
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+-----+-----+-----+-----+-----+-----+
|R|                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- R-Flag: Indicates the Root-ID length and if it is 16 octets value when set and is 4 octets value when clear.

- o Reserved: 2 octets which MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Root-ID: IPv4/IPv6 address of the head-end (Root) of the p2mp tree, based on AFI.
- o Tree-ID: a unique 4 octets identifier of the p2mp tree on the head-end (root)router.
- o Distinguisher: 4-octets value uniquely identifying the policy in the context of <Tree-ID, Originating Router's IP> tuple.

3.2. P2MP Policy Instance Descriptor

The P2MP Policy Instance Descriptor TLV identifies a P2MP Policy instance as defined in [draft-ietf-idr-sr-p2mp-policy]. This descriptor is used to report the state information of the P2MP Policy instance.

The TLV has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|               Type                 |               Length                 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Flags                 |               Reserved1                 |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Root-ID (4 or 16 octets)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Tree-ID (4 octets)                               |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Distinguisher (4 octets)                           |
+-----+-----+-----+-----+-----+-----+-----+-----+
//                               Node-ID (4 or 16 octets)                               //
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Instance-ID           |               Reserved2                 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

- o Type: TBD7
- o Length: variable

- o Flags: 2-octet field with following bit positions defined. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|R|N|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

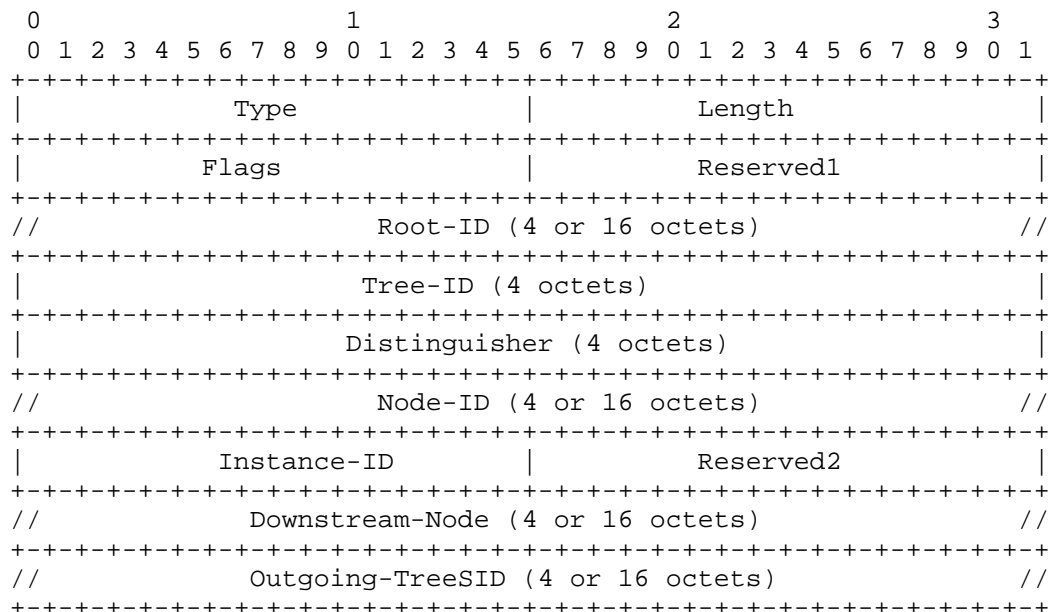
- R-Flag: Indicates the Root-ID length and if it is 16 octets value when set and is 4 octets value when clear.

- N-Flag: Indicates the Node-ID length and if it is 16 octets value when set and is 4 octets value when clear.

- o Reserved1: 2 octets which MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Root-ID: Same definition as in Section 3.1.
- o Tree-ID: Same definition as in Section 3.1.
- o Distinguisher: Same definition as in Section 3.1.
- o Node-ID: Advertising Node's IPv4/IPv6 address.
- o Instance-ID: 2-octets value uniquely identifying the path-instance with in the p2mp-policy.
- o Reserved2: 2-octets field which MUST be set to 0 by the originator and MUST be ignored by a receiver.


3.3. P2MP Policy downstream node Descriptor

The P2MP Policy downstream node Descriptor TLV identifies a downstream node in a P2MP Policy instance as defined in [draft-ietf-idr-sr-p2mp-policy].



where:

- o Type: TBD8
- o Length: variable
- o Flags: 2-octet field with following bit positions defined. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.



- R-Flag: Indicates the Root-ID length and if it is 16 octets value when set and is 4 octets value when clear.
- N-Flag: Indicates the Node-ID length and if it is 16 octets value when set and is 4 octets value when clear.
- D-Flag: Indicates the Downstream-Node length and if it is 16 octets value when set and is 4 octets value when clear.
- O-Flag: Indicates the Outgoing-TreeSID length and if it is 16 octets value when set and is 4 octets value when clear.

- o Reserved1: 2 octets which MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Root-ID: Same definition as in Section 3.1.
- o Tree-ID: Same definition as in Section 3.1.
- o Distinguisher: Same definition as in Section 3.1.
- o Node-ID: Advertising Node's IPv4/IPv6 address.
- o Instance-ID: 2-octets value uniquely identifying the path-instance with in the p2mp-policy.
- o Reserved2: 2-octets field which MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Downstream-Node: Downstream Node Identifier
- o Outgoing-TreeSID: The outgoing SID for this branch (MPLS or SRv6)

4. P2MP Policy States TLV

This section defines the various TLVs which enable the node to report the state of SR P2MP Policy. These TLVs (and their sub-TLVs) are carried in the optional non-transitive BGP Attribute "LINK_STATE Attribute" defined in [RFC7752] associated with the SR P2MP Policy NLRI type.

4.1. SR Policy Name TLV

This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Candidate Path Descriptor".

For the format definition of Segment Sub-TLV, refer to Section 5.4 of [draft-ietf-idr-bgp-ls-sr-policy-00].

4.2. SR Candidate Path Name TLV

This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Candidate Path Descriptor".

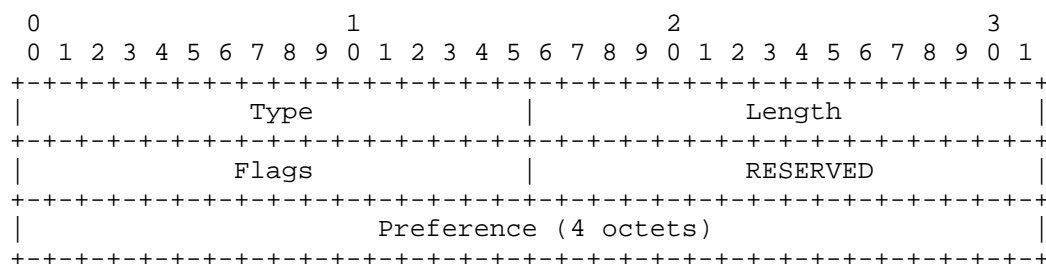
For the format definition of Segment Sub-TLV, refer to Section 5.5 of [draft-ietf-idr-bgp-ls-sr-policy-00].

4.3. SR P2MP Policy Candidate Path State TLV

The SR P2MP Policy Candidate Path (CP) State TLV provides the operational status and attributes of the SR P2MP Policy at the CP level. Only a single instance of this TLV is advertised for a given CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

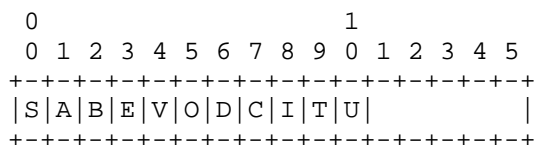
This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Candidate Path Descriptor".

The TLV has the following format:



where:

- o Type: TBD9
- o Length: 8 octets
- o Flags: 2-octet field that indicates attribute and status of the CP. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.



where:

- S-Flag: Indicates the CP is in an administrative shut state when set
- A-Flag: Indicates the CP is the active path (i.e. one provisioned in the forwarding plane) for the P2MP Policy when set

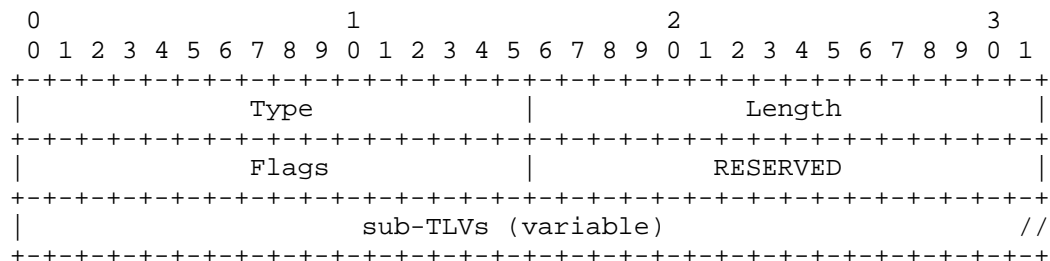
- B-Flag: Indicates the CP is the backup path (i.e. one identified for path protection of the active path) for the SR Policy when set
- E-Flag: Indicates that the CP has been evaluated for validity (e.g. headend may evaluate CPs based on their preferences) when set
- V-Flag: Indicates the CP has at least one valid instance when set. When the E-Flag is clear (i.e. the CP has not been evaluated), then this flag MUST be set to 0 by the originator and ignored by the receiver.
- D-Flag: Indicates the CP was delegated for computation to a PCE/controller when set.
- C-Flag: Indicates the CP was provisioned by a PCE/controller when set.
- I-Flag: Indicates the CP will perform the "drop upon invalid" behavior when no other active path is available for this P2MP Policy and this path is the one with the best preference amongst the available CPs.
- T-Flag: Indicates the CP has been marked as eligible for use as Transit Policy on the headend when set.
- U-Flag: Indicates the P2MP Policy that the CP belongs to is dropping traffic as a result of the "drop upon invalid" behavior being activated.
- o RESERVED: 2 octet2. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Preference: 4-octet value which indicates the preference of the CP.

4.4. SR P2MP Policy Leaf List TLV

The SR P2MP Policy leaf List TLV is used to report leaf list of a CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.


This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Candidate Path Descriptor".

The TLV has the following format:



where:

- o Type: TBD10
- o Length: variable
- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.



- D-Flag: Indicates the leaf list is comprised of SRv6 SIDs when set and indicates it is comprised of SR/MPLS labels when clear.

- V-Flag: Indicates the leaf list has passed verification or its verification was not required when set and failed verification when clear.

- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Sub-TLVs: variable and contains the ordered set of leaf and any other optional attributes associated with the specific leaf list.

4.4.1. SR P2MP Policy Leaf Sub-TLV

SR P2MP Policy Leaf Sub-TLV describes a single leaf in a leaf list. One or more instances of this sub-TLV in an ordered manner constitute a leaf list for a P2MP Policy candidate path.

The TLV has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |                                         |
|          Type                         |          Length                         |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |                                         |
|          Address (4 or 16 octets)    |                                         |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

- o Type: TBD11
- o Length: variable (valid values are 4, or 16)
- o Address: leaf's IPv4/IPv6 address

4.5. P2MP Policy Path Instance TLV

The P2MP Policy path instance TLV is used to report path instance of a P2MP Policy candidate path. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The TLV has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |                                         |
|          Type                         |          Length                         |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          Flags                       |          RESERVED                       |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         |                                         |
|          Sub-TLVs (variable)         |                                         |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

where:

- o Type: TBD12
- o Length: variable
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

[illegible]

- A-Flag: Indicates the instance-id is an active instance when set and the instance-id is not active instance when clear.

- o Sub-TLVs: variable and contains the ordered set of instance-id.

4.5.1. P2MP Policy Instance id Sub-TLV

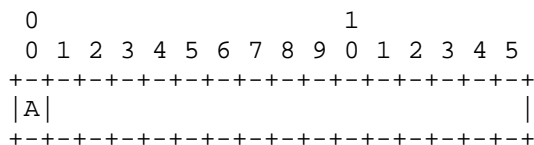
The P2MP Policy instance id sub-TLV describes a single instance in a path instance. One or more instances of this sub-TLV in an ordered manner constitute a leaf list for a P2MP Policy candidate path.

The TLV has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length																													
Flags										RESERVED																													
instance-id																				//																			

where:

- o Type: TBD13
- o Length: variable
- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.



- A-Flag: Indicates the instance-id is an active instance when set and the instance-id is not active instance when clear.

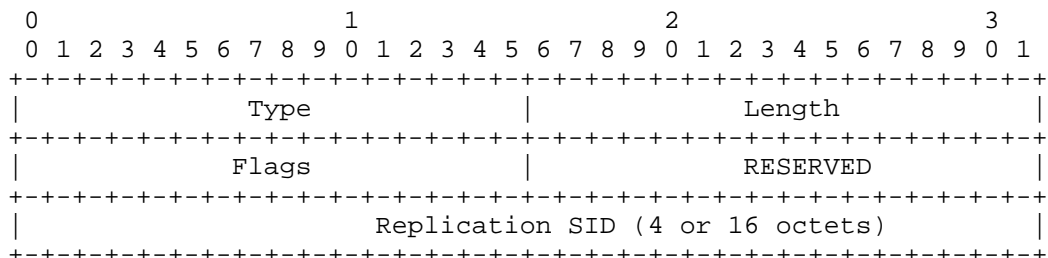
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Instance-id: 4-octect field identifier.

4.6. P2MP Policy Replication Segment TLV

The P2MP Policy replication segment TLV describes the incoming replication segment of a P2MP Policy Instance. Only a single instance of this TLV is advertised for a given instance. If multiple instances are present, then the first one is considered valid and the rest are ignored.

This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Instance Descriptor".

The TLV has the following format:



where:

- o Type: TBD
- o Length: variable(valid value is 8 or 20)
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+---+---+---+---+---+---+---+---+---+
|D|B|U|F|                               |
+---+---+---+---+---+---+---+---+---+

```

- D-Flag: Indicates the dataplane for the "Replication SID" and if they are 16 octet SRv6 SID when set and are 4 octet SR/MPLS label value when clear.

- B-Flag: Indicates the allocation of the value in the "Replication SID" field when set and indicates that BSID is not allocated when clear.

- U-Flag: Indicates the specified "Replication SID" value is unavailable when set.

- F-Flag: Indicates the BSID value is one allocated dynamically due to fallback (e.g. when specified "Replication SID" is unavailable) when set.

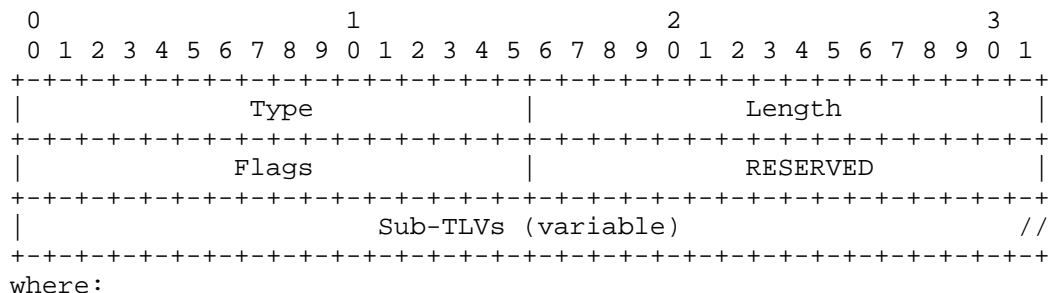
- o Replication SID: It indicates the incoming replication SID of instance.

4.7. P2MP Policy Downstream List TLV


The P2MP Policy downstream list TLV describes the downstream node list of a P2MP Policy Instance. Only a single instance of this TLV is advertised for a given instance. If multiple instances are present, then the first one is considered valid and the rest are ignored.

This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Instance Descriptor".

The TLV has the following format:



- o Type: TBD
- o Length: variable
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.



- D-Flag: Indicates the dataplane for the "Replication SID" and if they are 16 octet SRv6 SID when set and are 4 octet SR/MPLS label value when clear.

- B-Flag: Indicates the allocation of the value in the "Replication SID" field when set and indicates that BSID is not allocated when clear.

- U-Flag: Indicates the specified "Replication SID" value is unavailable when set.

- F-Flag: Indicates the BSID value is one allocated dynamically due to fallback (e.g. when specified "Replication SID" is unavailable) when set.

- o Sub-TLVs: variable and contains the ordered set of downstream node and any other optional attributes associated with the specific downstream list.

4.7.1. P2MP Policy Downstream node Sub-TLV

The P2MP Policy downstream node Sub-TLV describes a single downstream node in a downstream list. One or more instances of this sub-TLV in an ordered manner constitute a downstream list for a P2MP Policy instance.

This TLV is used when the root node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy Instance Descriptor".

The TLV has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|               Type                 |               Length                 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Flags                 |               RESERVED                 |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|               Downstream-Node       |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|               Outgoing-TreeSID      |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|               Sub-TLVs (variable)   |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
where:

```

- o Type: TBD
- o Length: variable
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Flags: 2-octet field that indicates attribute and status of the downstream node. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

```

      0               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+-----+-----+-----+-----+-----+-----+-----+
|D|O|                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```


- D-Flag: Indicates the downstream-Node length and if it is 16 octets value when set and is 4 octets value when clear.

- O-Flag: Indicates the Outgoing-TreeSID length and if it is 16 octets value when set and is 4 octets value when clear.
- o Downstream-Node: Downstream Node Identifier
- o Outgoing-TreeSID: The outgoing SID for this branch (MPLS or SRv6)
- o Sub-TLVs: variable and contains the ordered set of optional attributes associated with the specific downstream node.

4.8. P2MP Policy Downstream-node Segment list TLV

The P2MP Policy Downstream Segment list TLV describes a segment list for a downstream node. One or more instances of this TLV in an ordered manner constitute a set of segment list for a downstream node.

This TLV is used when the advertising node distributing the NLRI whose P2MP Policy Descriptor is "P2MP Policy downstream node Descriptor ".



where:

- o Type: TBD
- o Length: variable
- o Flags: 2-octet field that indicates attribute and status of the leaf list. The following is the bit positions definition. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

```

      0                               1
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-----+
|D|E|C|V|R|F|A|M|T|          |
+-----+

```

- D-Flag: Indicates the SID-List is comprised of SRv6 SIDs when set and indicates it is comprised of SR/MPLS labels when clear.
- E-Flag: Indicates that SID-List is an explicit path when set and indicates a dynamic path when clear.
- C-Flag: Indicates that SID-List has been computed for a dynamic path when set. It is always reported as set for explicit paths.
- V-Flag: Indicates the SID-List has passed verification or its verification was not required when set and failed verification when clear.
- R-Flag: Indicates that the first Segment has been resolved when set and failed resolution when clear.
- F-Flag: Indicates that the computation for the dynamic path failed when set and succeeded (or not required in case of explicit path) when clear
- A-Flag: Indicates that all the SIDs in the SID-List belong to the specified algorithm when set.
- T-Flag: Indicates that all the SIDs in the SID-List belong to the specified topology (identified by the multi-topology ID) when set.
- M-Flag: Indicates that the SID-list has been removed from the forwarding plane due to fault detection by a monitoring mechanism (e.g. BFD) when set and indicates no fault detected or monitoring is not being done when clear.
- o RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- o Segment List id: 2 octets value indicating the segment list id
- o Protection Segment List id: 2 octets value indicating the segment list id that is being used as protection. If there is no segment for protection, this field should be set 0.
- o Weight: 4 octets value indicating the weight associated with segment list.

- o Sub-TLVs: variable and contains the ordered set of segment any other optional attributes associated with the specific segment list.

4.8.1. Segment Sub-TLV

The segment Sub-TLV describes a single segment in a segment list. One or more instances of this sub-TLV in an ordered manner constitute a segment list for a P2MP Policy downstream-node.

For the format definition of Segment Sub-TLV, refer to Section 5.8 of [draft-ietf-idr-bgp-ls-sr-policy-00].

5. IANA Considerations

5.1. BGP-LS NLRI-Type

IANA maintains a registry called "BGP-LS NLRI-Types" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

This document is to apply for a P2MP Policy NLRI-Type.

Type	NLRI Type	Reference
TBD	P2MP Policy NLRI	this document

5.2. BGP-LS Protocol-ID

IANA maintains a registry called "BGP-LS Protocol-IDs" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

This document is to apply for a P2MP Policy protocol-ID.

Protocol-ID	NLRI information source protocol	Reference
TBD	P2MP Policy	this document

5.3. BGP-LS TLV

IANA maintains a registry called "Node Anchor, Link Descriptor and Link Attribute TLVs" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

This document is to apply for a set of TLVs

Code Point	Description	Value defined in
TBD	P2MP Policy Candidate Path Descriptor	this document
TBD	P2MP Policy Instance Descriptor	this document
TBD	P2MP Policy downstream node Descriptor	this document
TBD	P2MP Policy Candidate Path State	this document
TBD	P2MP Policy leaf List	this document
TBD	P2MP Policy leaf	this document
TBD	P2MP Policy Path Instance	this document
TBD	P2MP Policy Instance id	this document
TBD	P2MP Policy Replication Segment	this document
TBD	P2MP Policy Downstream List	this document
TBD	P2MP Policy Downstream Node	this document
TBD	P2MP Policy Downstream Node segment list	this document

6. Security Considerations

TBD

7. References

7.1. 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>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.

- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [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>>.
- [RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August 2021, <<https://www.rfc-editor.org/info/rfc9086>>.
- [draft-ietf-pim-sr-p2mp-policy-06] D. Voyer, C. Filsfils, H. Bidgoli, Z. Zhang, "Segment Routing Point-to-Multipoint Policy", draft-ietf-pim-sr-p2mp-policy-06, DOI 10.17487/draft-ietf-pim-sr-p2mp-policy-06, April 2023.
- [draft-ietf-idr-sr-p2mp-policy-00] Hooman Bidgoli, Daniel Voyer, Andrew Stone, Rishabh Parekh, Serge Krier, Swadesh Agrawal, "Advertising p2mp policies in BGP", draft-ietf-idr-sr-p2mp-policy-00, DOI 10.17487/draft-ietf-idr-sr-p2mp-policy-00, May 2022.
- [draft-ietf-idr-bgp-ls-sr-policy-00] S. Previdi, K. Talaulikar, J. Dong, H. Gredler, J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", draft-ietf-idr-bgp-ls-sr-policy-00, DOI 10.17487/draft-ietf-idr-bgp-ls-sr-policy-00, March 2023.

Contributors

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Changwang Lin
New H3C Technologies
Beijing
China

Email: linchangwang.04414@h3c.com

Hooman Bidgoli
Nokia
Ottawa
Canada
Email: hooman.bidgoli@nokia.com

Zheng(Sandy) Zhang
ZTE Corporation
Email: zhang.zheng@zte.com.cn

