

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 22, 2017

S. Previdi, Ed.
C. Filsfils
Cisco Systems, Inc.
K. Patel
Arrcus, Inc.
S. Ray
Individual Contributor
J. Dong
Huawei Technologies
June 20, 2017

**BGP-LS extensions for Segment Routing BGP Egress Peer Engineering
draft-ietf-idr-bgpls-segment-routing-epe-13**

Abstract

Segment Routing (SR) leverages source routing. A node steers a packet through a controlled set of instructions, called segments, by prepending the packet with an SR header. A segment can represent any instruction, topological or service-based. SR allows to enforce a flow through any topological path and service chain while maintaining per-flow state only at the ingress node of the SR domain.

The Segment Routing architecture can be directly applied to the MPLS dataplane with no change on the forwarding plane. It requires minor extension to the existing link-state routing protocols.

This document outline a BGP-LS extension for exporting BGP peering node topology information (including its peers, interfaces and peering ASs) in a way that is exploitable in order to compute efficient BGP Peering Engineering policies and strategies.

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 <http://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 December 22, 2017.

Copyright Notice

Copyright (c) 2017 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
2.	Segment Routing Documents	3
3.	BGP Peering Segments	4
4.	Link NLRI for BGP-EPE Connectivity Description	5
4.1.	BGP Router ID and Member ASN	5
4.2.	Mandatory BGP-EPE Node Descriptors	6
4.3.	Optional BGP-EPE Node Descriptors	7
4.4.	Link Attributes	7
5.	Peer-Node and Peer-Adj SIDs	9
5.1.	Peer-Node-SID	10
5.2.	Peer-Adj-SID	11
5.3.	Peer-Set-SID	12
6.	Illustration	12
6.1.	Reference Diagram	12
6.2.	Peer-Node-SID for Node D	14
6.3.	Peer-Node-SID for Node F	15
6.4.	Peer-Node-SID for Node E	15
6.5.	Peer-Adj-SID for Node E, Link 1	15
6.6.	Peer-Adj-SID for Node E, Link 2	16
7.	Implementation Status	16
8.	IANA Considerations	17
8.1.	New BGP-LS Protocol-ID	17
8.2.	Node Descriptors and Link Attribute TLVs	18
9.	Manageability Considerations	18

10.	Security Considerations	19
11.	Contributors	19
12.	Acknowledgements	19
13.	References	19
13.1.	Normative References	19
13.2.	Informative References	20
	Authors' Addresses	21

[1.](#) Introduction

Segment Routing (SR) leverages source routing. A node steers a packet through a controlled set of instructions, called segments, by prepending the packet with an SR header with segment identifiers (SID). A SID can represent any instruction, topological or service-based. SR allows to enforce a flow through any topological path and service chain while maintaining per-flow state only at the ingress node of the SR domain.

The Segment Routing architecture can be directly applied to the MPLS dataplane with no change on the forwarding plane. It requires minor extension to the existing link-state routing protocols.

This document outline a BGP-LS extension for exporting BGP peering node topology information (including its peers, interfaces and peering ASs) in a way that is exploitable in order to compute efficient BGP Egress Peer Engineering (BGP-EPE) policies and strategies.

This document defines the BGP-LS extensions required to support the Peer Node SID describing the BGP session between two nodes, the Peer Adjacency SID describing the link (one or more) that is used by the BGP session and the Peer Set SID describing an arbitrary set of sessions or links between the local BGP node and its peers. These SIDs represent the segments defined in [[I-D.ietf-spring-segment-routing-central-epe](#)].

While an egress point topology usually refers to eBGP sessions between external peers, there's nothing in the extensions defined in this document that would prevent the use of these extensions in the context of iBGP sessions.

[2.](#) Segment Routing Documents

The main reference for this document is the SR architecture defined in [[I-D.ietf-spring-segment-routing](#)].

The Segment Routing BGP Egress Peer Engineering (BGP-EPE) architecture is described in [\[I-D.ietf-spring-segment-routing-central-epe\]](#).

3. BGP Peering Segments

As defined in [\[I-D.ietf-spring-segment-routing-central-epe\]](#), a BGP-EPE enabled Egress PE node MAY advertise SIDs corresponding to its attached peers. These SIDs are called BGP peering segments or BGP Peering SIDs. In case of eBGP, they enable the expression of source-routed inter-domain paths.

An ingress border router of an AS may compose a list of SIDs to steer a flow along a selected path within the AS, towards a selected egress border router C of the AS and through a specific peer. At minimum, a BGP-EPE policy applied at an ingress PE involves two SIDs: the Node SID of the chosen egress PE and then the BGP Peering SID for the chosen egress PE peer or peering interface.

This document defines the BGP-LS extensions for the BGP-EPE Peering SIDs:

- o Peer Node Segment (Peer-Node-SID)
- o Peer Adjacency Segment (Peer-Adj-SID)
- o Peer Set Segment (Peer-Set-SID)

that have been defined in [\[I-D.ietf-spring-segment-routing-central-epe\]](#).

Each BGP session MUST be described by a Peer Node SID. The description of the BGP session MAY be augmented by additional Adjacency SIDs. Finally, each Peer Node SID and Peer Adjacency SID MAY be part of the same group/set so to be able to group EPE resources under a common Peer-Set SID.

Therefore, when the extensions defined in this document are applied to the use case defined in [\[I-D.ietf-spring-segment-routing-central-epe\]](#):

- o One Peer-Node-SID MUST be present.
- o One or more Peer-Adj-SID MAY be present.
- o Each of the Peer-Node-SID Peer-Adj-SID MAY use the same Peer-Set-SID.

While an egress point topology usually refers to eBGP sessions between external peers, there's nothing in the extensions defined in this document that would prevent the use of these extensions in the context of iBGP sessions.

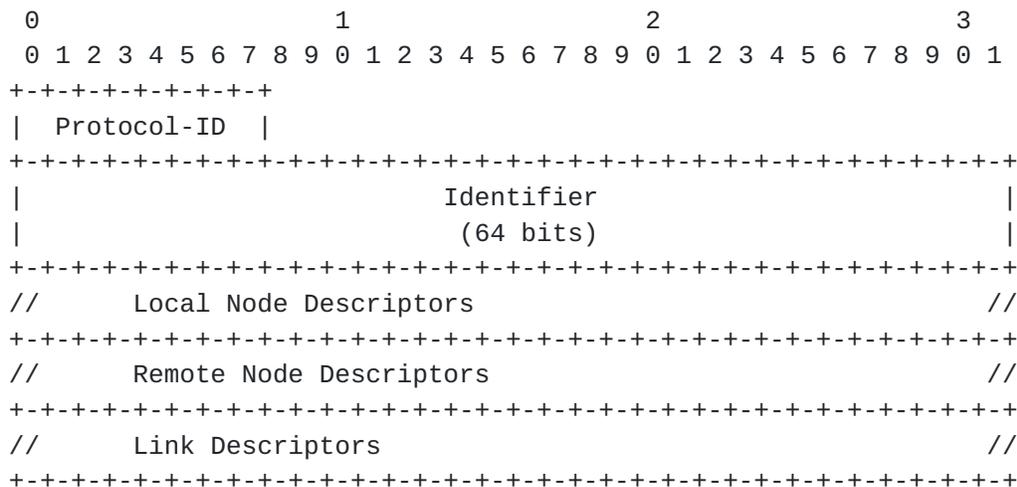
4. Link NLRI for BGP-EPE Connectivity Description

This section describes the NLRI used for describing the connectivity of the BGP Egress router. The connectivity is based on links and remote peers/ASs and therefore the existing Link NLRI Type (defined in [RFC7752]) is used. A new Protocol-ID is used: BGP (codepoint 7 assigned by IANA (Section 8) from the registry "BGP-LS Protocol-IDs").

The use of a new Protocol-ID allows separation and differentiation between the NLRIs carrying BGP-EPE descriptors from the NLRIs carrying IGP link-state information as defined in [RFC7752]. The Link NLRI Type uses descriptors and attributes already defined in [RFC7752] in addition to new TLVs defined in the following sections of this document.

The extensions defined in this document apply to both internal and external BGP-LS EPE advertisements.

[RFC7752] defines Link NLRI Type is as follows:



Node Descriptors and Link Descriptors are defined in [RFC7752].

4.1. BGP Router ID and Member ASN

Two new Node Descriptors Sub-TLVs are defined in this document:

- o BGP Router Identifier (BGP Router-ID):

Type: 516 (assigned by IANA ([Section 8](#)) from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs").

Length: 4 octets

Value: 4 octet unsigned integer representing the BGP Identifier as defined in [[RFC4271](#)] and [[RFC6286](#)].

o Confederation Member ASN (Member-ASN)

Type: 517 (assigned by IANA ([Section 8](#)) from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs").

Length: 4 octets

Value: 4 octet unsigned integer representing the Member ASN inside the Confederation. [[RFC5065](#)].

[4.2.](#) Mandatory BGP-EPE Node Descriptors

The following Node Descriptors Sub-TLVs MUST appear in the Link NLRI as Local Node Descriptors:

- o BGP Router-ID, which contains the BGP Identifier of the local BGP-EPE capable node.
- o Autonomous System Number, which contains the local ASN or local confederation identifier (ASN) if confederations are used.
- o BGP-LS Identifier.

It has to be noted that [[RFC6286](#)] ([section 2.1](#)) requires the BGP identifier (router-id) to be unique within an Autonomous System. Therefore, the <ASN, BGP identifier> tuple is globally unique.

The following Node Descriptors Sub-TLVs MUST appear in the Link NLRI as Remote Node Descriptors:

- o BGP Router-ID, which contains the BGP Identifier of the peer node.
- o Autonomous System Number, which contains the peer ASN or the peer confederation identifier (ASN), if confederations are used.

4.3. Optional BGP-EPE Node Descriptors

The following Node Descriptors Sub-TLVs MAY appear in the Link NLRI as Local Node Descriptors:

- o Member-ASN, which contains the ASN of the confederation member (when BGP confederations are used).
- o Node Descriptors as defined in [[RFC7752](#)].

The following Node Descriptors Sub-TLVs MAY appear in the Link NLRI as Remote Node Descriptors:

- o Member-ASN, which contains the ASN of the confederation member (when BGP confederations are used).
- o Node Descriptors as defined in defined in [[RFC7752](#)].

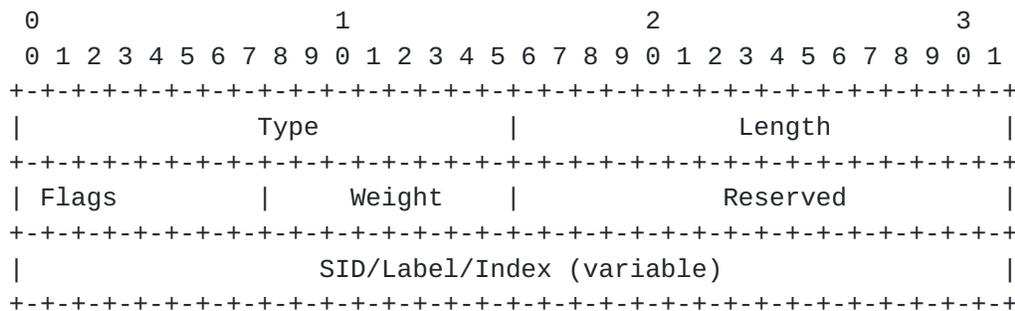
4.4. Link Attributes

The following BGP-LS Link attributes TLVs are used with the Link NLRI:

TLV Code	Description	Length
1101	Peer Node Segment Identifier (Peer-Node-SID)	variable
1102	Peer Adjacency Segment Identifier (Peer-Adj-SID)	variable
1103	Peer Set Segment Identifier (Peer-Set-SID)	variable

Figure 1: BGP-LS TLV code points for BGP-EPE

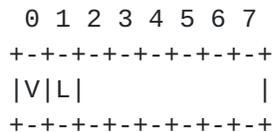
Peer-Node-SID, Peer-Adj-SID and Peer-Set-SID have all the same format defined here below:



where:

Figure 2

- o Type: 1101 or 1102 or 1103 (assigned by IANA ([Section 8](#)) from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs").
- o Length: variable.
- o Flags: one octet of flags used when advertising a Peer-Adj-SID (Peer-Node and Peer-Set SIDs don't have flags defined). The following Peer-Adj-SID flags have been defined:



where:

- * V-Flag: Value flag. If set, then the Adj-SID carries a value. By default the flag is SET.
- * L-Flag: Local Flag. If set, then the value/index carried by the Adj-SID has local significance. By default the flag is SET.
- * Other bits: MUST be zero when originated and ignored when received.

- o Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing. An example use of the weight is described in [[I-D.ietf-spring-segment-routing](#)].
- o SID/Index/Label. According to the TLV length and to the V and L flags settings, it contains either:

- * A 3 octet local label where the 20 rightmost bits are used for encoding the label value. In this case the V and L flags MUST be set.
- * A 4 octet index defining the offset in the SRGB (Segment Routing Global Block as defined in [\[I-D.ietf-spring-segment-routing\]](#) advertised by this router. In this case, the SRGB MUST be advertised using the extensions defined in [\[I-D.ietf-idr-bgp-ls-segment-routing-ext\]](#).
- * A 16 octet IPv6 address. In this case the V flag MUST be set. The L flag MUST be unset if the IPv6 address is globally unique.

The values of the Peer-Node-SID, Peer-Adj-SID and Peer-Set-SID Sub-TLVs SHOULD be persistent across router restart.

The Peer-Node-SID MUST be present when BGP-LS is used for the use case described in [\[I-D.ietf-spring-segment-routing-central-epe\]](#) and MAY be omitted for other use cases.

The Peer-Adj-SID and Peer-Set-SID SubTLVs MAY be present when BGP-LS is used for the use case described in [\[I-D.ietf-spring-segment-routing-central-epe\]](#) and MAY be omitted for other use cases.

In addition, BGP-LS Nodes and Link Attributes, as defined in [\[RFC7752\]](#) MAY be inserted in order to advertise the characteristics of the link.

5. Peer-Node and Peer-Adj SIDs

In this section the following SIDs are defined:

Peer Node Segment Identifier (Peer-Node-SID)

Peer Adjacency Segment Identifier (Peer-Adj-SID)

Peer Set Segment Identifier (Peer-Set-SID)

The Peer-Node, Peer-Adj and Peer-Set SIDs can be either a local or a global (depending on the setting of the V and L flags defined in Figure 2. For example, when BGP-EPE is used in the context of a SR network over the IPv6 dataplane, it is likely the case that the IPv6 addresses used as SIDs will be global.

5.1. Peer-Node-SID

The Peer-Node-SID describes the BGP session peer (neighbor). It MUST be present when describing a BGP-EPE topology as defined in [\[I-D.ietf-spring-segment-routing-central-epe\]](#). The Peer-Node-SID is encoded within the BGP-LS Link NLRI specified in [Section 4](#).

The Peer-Node-SID, at the BGP node advertising it, has the following semantic:

- o SR header operation: NEXT (as defined in [\[I-D.ietf-spring-segment-routing\]](#)).
- o Next-Hop: the connected peering node to which the segment is related.

The Peer-Node-SID is advertised with a Link NLRI, where:

- o Local Node Descriptors contains
 - Local BGP Router-ID of the BGP-EPE enabled egress PE.
 - Local ASN.
 - BGP-LS Identifier.
- o Remote Node Descriptors contains
 - Peer BGP Router-ID (i.e.: the peer BGP ID used in the BGP session).
 - Peer ASN.
- o Link Descriptors Sub-TLVs, as defined in [\[RFC7752\]](#), contain the addresses used by the BGP session:
 - * IPv4 Interface Address (Sub-TLV 259) contains the BGP session IPv4 local address.
 - * IPv4 Neighbor Address (Sub-TLV 260) contains the BGP session IPv4 peer address.
 - * IPv6 Interface Address (Sub-TLV 261) contains the BGP session IPv6 local address.
 - * IPv6 Neighbor Address (Sub-TLV 262) contains the BGP session IPv6 peer address.
- o Link Attribute contains the Peer-Node-SID TLV as defined in [Section 4.4](#).

- o In addition, BGP-LS Link Attributes, as defined in [[RFC7752](#)], MAY be inserted in order to advertise the characteristics of the link.

5.2. Peer-Adj-SID

The Peer-Adj-SID, at the BGP node advertising it, has the following semantic:

- o SR header operation: NEXT (as defined in [[I-D.ietf-spring-segment-routing](#)]).
- o Next-Hop: the interface peer address.

The Peer-Adj-SID is advertised with a Link NLRI, where:

- o Local Node Descriptors contains

Local BGP Router-ID of the BGP-EPE enabled egress PE.

Local ASN.

BGP-LS Identifier.

- o Remote Node Descriptors contains

Peer BGP Router-ID (i.e.: the peer BGP ID used in the BGP session).

Peer ASN.

- o Link Descriptors Sub-TLVs, as defined in [[RFC7752](#)], MUST contain the following TLVs:

- * Link Local/Remote Identifiers (Sub-TLV 258) contains the 4-octet Link Local Identifier followed by the 4-octet value 0 indicating the Link Remote Identifier in unknown [[RFC5307](#)].

- o In addition, Link Descriptors Sub-TLVs, as defined in [[RFC7752](#)], MAY contain the following TLVs:

- * IPv4 Interface Address (Sub-TLV 259) contains the address of the local interface through which the BGP session is established.
- * IPv6 Interface Address (Sub-TLV 261) contains the address of the local interface through which the BGP session is established.
- * IPv4 Neighbor Address (Sub-TLV 260) contains the IPv4 address of the peer interface used by the BGP session.

- * IPv6 Neighbor Address (Sub-TLV 262) contains the IPv6 address of the peer interface used by the BGP session.
- o Link attribute used with the Peer-Adj-SID contains the TLV as defined in [Section 4.4](#).

In addition, BGP-LS Link Attributes, as defined in [[RFC7752](#)], MAY be inserted in order to advertise the characteristics of the link.

5.3. Peer-Set-SID

The Peer-Set-SID, at the BGP node advertising it, has the following semantic:

- o SR header operation: NEXT (as defined in [[I-D.ietf-spring-segment-routing](#)]).
- o Next-Hop: load balance across any connected interface to any peer in the related set.

The Peer-Set-SID is advertised within a Link NLRI (describing a Peer Node Segment or a Peer Adjacency segment) as a BGP-LS attribute.

The Peer Set Attribute contains the Peer-Set-SID TLV, defined in [Section 4.4](#) identifying the set of which the Peer-Node-SID or Peer-Adj-SID is a member.

6. Illustration

6.1. Reference Diagram

The following reference diagram is used throughout this document. The solution is illustrated for IPv6 with MPLS-based SIDs and the BGP-EPE topology is based on eBGP sessions between external peers.

As stated in [Section 3](#), the solution illustrated hereafter is equally applicable to an iBGP session topology. In other words, the solution also applies to the case where C, D, F, and E are in the same AS and run iBGP sessions between each other.

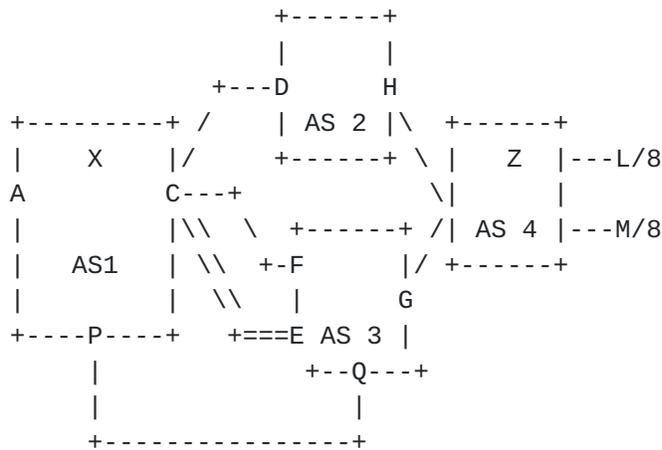


Figure 3: Reference Diagram

IP addressing:

- o C's IP address of interface to D: 2001:db8:cd::c/64, D's interface: 2001:db8:cd::d/64
- o C's IP address of interface to F: 2001:db8:cf::c/64, F's interface: 2001:db8:cf::f/64
- o C's IP address of upper interface to E: 2001:db8:ce1::c/64, E's interface: 2001:db8:ce1::e
- o C's local identifier of upper interface to E: 0.0.0.1.0.0.0.0
- o C's IP address of lower interface to E: 2001:db8:ce2::c, E's interface: 2001:db8:ce2::e
- o C's local identifier of lower interface to E: 0.0.0.2.0.0.0.0
- o Loopback of E used for eBGP multi-hop peering to C: 2001:db8:e::e/128
- o C's loopback is 2001:db8:c::c/128 with SID 64

BGP Router-IDs are C, D, F and E.

- o C's BGP Router-ID: 192.0.2.3
- o D's BGP Router-ID: 192.0.2.4
- o E's BGP Router-ID: 192.0.2.5
- o F's BGP Router-ID: 192.0.2.6

C's BGP peering:

- o Single-hop eBGP peering with neighbor 2001:db8:cd::d (D)
- o Single-hop eBGP peering with neighbor 2001:db8:cf::f (F)
- o Multi-hop eBGP peering with E on ip address 2001:db8:e::e (E)

C's resolution of the multi-hop eBGP session to E:

- o Static route 2001:db8:e::e/128 via 2001:db8:ce1::e
- o Static route 2001:db8:e::e/128 via 2001:db8:ce2::e

Node C configuration is such that:

- o A Peer-Node-SID is allocated to each peer (D, F and E).
- o An Peer-Adj-SID is defined for each recursing interface to a multi-hop peer (CE upper and lower interfaces).
- o A Peer-Set-SID is defined to include all peers in AS3 (peers F and E).

Local BGP-LS Identifier in router C is set to 10000.

The Link NLRI Type is used in order to encode C's connectivity. The Link NLRI uses the Protocol-ID value (to be assigned by IANA)

Once the BGP-LS update is originated by C, it may be advertised to internal (iBGP) as well as external (eBGP) neighbors supporting the BGP-LS EPE extensions defined in this document.

6.2. Peer-Node-SID for Node D

Descriptors:

- o Local Node Descriptors (BGP Router-ID, local ASN, BGP-LS Identifier): 192.0.2.3, AS1, 10000
- o Remote Node Descriptors (BGP Router-ID, peer ASN): 192.0.2.4, AS2
- o Link Descriptors (BGP session IPv6 local address, BGP session IPv6 neighbor address): 2001:db8:cd::c, 2001:db8:cd::d

Attributes:

- o Peer-Node-SID: 1012

- o Link Attributes: see [section 3.3.2 of \[RFC7752\]](#)

6.3. Peer-Node-SID for Node F

Descriptors:

- o Local Node Descriptors (BGP Router-ID, ASN, BGPLS Identifier): 192.0.2.3, AS1, 10000
- o Remote Node Descriptors (BGP Router-ID ASN): 192.0.2.6, AS3
- o Link Descriptors (BGP session IPv6 local address, BGP session IPv6 peer address): 2001:db8:cf::c, 2001:db8:cf::f

Attributes:

- o Peer-Node-SID: 1022
- o Peer-Set-SID: 1060
- o Link Attributes: see [section 3.3.2 of \[RFC7752\]](#)

6.4. Peer-Node-SID for Node E

Descriptors:

- o Local Node Descriptors (BGP Router-ID, ASN, BGP-LS Identifier): 192.0.2.3, AS1, 10000
- o Remote Node Descriptors (BGP Router-ID, ASN): 192.0.2.5, AS3
- o Link Descriptors (BGP session IPv6 local address, BGP session IPv6 peer address): 2001:db8:c::c, 2001:db8:e::e

Attributes:

- o Peer-Node-SID: 1052
- o Peer-Set-SID: 1060

6.5. Peer-Adj-SID for Node E, Link 1

Descriptors:

- o Local Node Descriptors (BGP Router-ID, ASN, BGP-LS Identifier): 192.0.2.3, AS1, 10000
- o Remote Node Descriptors (BGP Router-ID, ASN): 192.0.2.5, AS3

- o Link Descriptors (local interface identifier, IPv6 peer interface address): 0.0.0.1.0.0.0.0 , 2001:db8:ce1::e

Attributes:

- o Peer-Adj-SID: 1032
- o LinkAttributes: see [section 3.3.2 of \[RFC7752\]](#)

6.6. Peer-Adj-SID for Node E, Link 2

Descriptors:

- o Local Node Descriptors (BGP Router-ID, ASN, BGP-LS Identifier): 192.0.2.3, AS1, 10000
- o Remote Node Descriptors (BGP Router-ID, ASN): 192.0.2.5, AS3
- o Link Descriptors (local interface identifier, IPv6 peer interface address): 0.0.0.2.0.0.0.0 , 2001:db8:ce2::e

Attributes:

- o Peer-Adj-SID: 1042
- o LinkAttributes: see [section 3.3.2 of \[RFC7752\]](#)

7. Implementation Status

Note to RFC Editor: Please remove this section prior to publication, as well as the reference to [RFC 7942](#).

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [\[RFC7942\]](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [\[RFC7942\]](#), "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation

and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

Several early implementations exist and will be reported in detail in a forthcoming version of this document. For purposes of early interoperability testing, when no FCFS code point was available, implementations have made use of the following values:

Codepoint	Description
7	Protocol-ID BGP
516	BGP Router-ID
517	BGP Confederation Member
1101	Peer-Node-SID
1102	Peer-Adj-SID
1103	Peer-Set-SID

IANA has now confirmed the assignment of the above coidepoints. See [Section 8](#).

8. IANA Considerations

This document defines:

A new Protocol-ID: BGP. The codepoint is from the "BGP-LS Protocol-IDs" registry.

Two new TLVs: BGP-Router-ID and BGP Confederation Member. The codepoints are in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

Three new BGP-LS Attribute TLVs: Peer-Node-SID, Peer-Adj-SID and Peer-Set-SID. The codepoints are in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

8.1. New BGP-LS Protocol-ID

This document defines a new value in the registry "BGP-LS Protocol-IDs":

Codepoint	Description	Status
7	BGP	Assigned by IANA

8.2. Node Descriptors and Link Attribute TLVs

This document defines 5 new TLVs in the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs":

- o Two new node descriptor TLVs
- o Three new link attribute TLVs

All the new 5 codepoints are in the same registry: "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs". However, the registry is organized in ranges (node descriptors, link descriptors, node attributes, link attributes).

The following new Node Descriptors TLVs are defined:

```

+-----+
| Codepoint | Description          | Status          |
+-----+
|    516    | BGP Router-ID       | Assigned by IANA |
|    517    | BGP Confederation Member | Assigned by IANA |
+-----+

```

The following new Link Attribute TLVs are defined:

```

+-----+
| Codepoint | Description          | Status          |
+-----+
|    1101   | Peer-Node-SID       | Assigned by IANA |
|    1102   | Peer-Adj-SID        | Assigned by IANA |
|    1103   | Peer-Set-SID        | Assigned by IANA |
+-----+

```

9. Manageability Considerations

The BGP-LS ([\[RFC7752\]](#)) extensions that are described in this document consist of additional BGP-LS descriptors and TLVs that will follow the same manageability functions of BGP-LS, described in [\[RFC7752\]](#).

The operator MUST be capable of configuring, enabling, disabling the advertisement of each of the Peer-Node-SID, Peer-Adj-SID and Peer-Set-SID as well as to control which information is advertised to which internal or external peer. This is not different from what is required by a BGP speaker in terms of information origination and advertisement. In addition, the advertisement of EPE information MUST conform to standard BGP advertisement and propagation rules (iBGP, eBGP, Route-Reflectors, Confederations).

10. Security Considerations

[RFC7752] defines BGP-LS NLRIs to which the extensions defined in this document apply.

The Security Section of [[RFC7752](#)] also applies to:

- o New Node Descriptors Sub-TLVs: BGP-Router-ID and BGP-Confederation-Member;
- o New BGP-LS Attributes TLVs: Peer-Node-SID, Peer-Adj-SID and Peer-Set-SID.

The extensions defined in this document do not introduce any additional security aspects of BGP-LS.

11. Contributors

Mach (Guoyi) Chen
Huawei Technologies
China

Email: mach.chen@huawei.com

Acee Lindem
Cisco Systems Inc.
US

Email: acee@cisco.com

12. Acknowledgements

The authors would like to thank Jakob Heitz, Howard Yang, Hannes Gredler, Peter Psenak, Ketan Jivan Talaulikar, Arjun Sreekantiah and Bruno Decraene for their feedback and comments.

13. References

13.1. Normative References

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S.,
and R. Shakir, "Segment Routing Architecture", [draft-ietf-spring-segment-routing-11](#) (work in progress), February 2017.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<http://www.rfc-editor.org/info/rfc4271>>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", [RFC 5065](#), DOI 10.17487/RFC5065, August 2007, <<http://www.rfc-editor.org/info/rfc5065>>.
- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 5307](#), DOI 10.17487/RFC5307, October 2008, <<http://www.rfc-editor.org/info/rfc5307>>.
- [RFC6286] Chen, E. and J. Yuan, "Autonomous-System-Wide Unique BGP Identifier for BGP-4", [RFC 6286](#), DOI 10.17487/RFC6286, June 2011, <<http://www.rfc-editor.org/info/rfc6286>>.

13.2. Informative References

- [I-D.ietf-idr-bgp-ls-segment-routing-ext]
Previdi, S., Psenak, P., Filsfils, C., Gredler, H., Chen, M., and j. jeffrant@gmail.com, "BGP Link-State extensions for Segment Routing", [draft-ietf-idr-bgp-ls-segment-routing-ext-01](#) (work in progress), February 2017.
- [I-D.ietf-spring-segment-routing-central-epe]
Filsfils, C., Previdi, S., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", [draft-ietf-spring-segment-routing-central-epe-05](#) (work in progress), March 2017.
- [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, <<http://www.rfc-editor.org/info/rfc7752>>.
- [RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [BCP 205](#), [RFC 7942](#), DOI 10.17487/RFC7942, July 2016, <<http://www.rfc-editor.org/info/rfc7942>>.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems, Inc.

Email: stefano@previdi.net

Clarence Filsfils
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com

Keyur Patel
Arrcus, Inc.

Email: Keyur@arrcus.com

Saikat Ray
Individual Contributor

Email: raysaikat@gmail.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

