Inter-Domain Routing   S. Previdi
Internet-Draft   Individual
Intended status: Standards Track   K. Talaulikar, Ed.
Expires: September 25, 2019   C. Filsfils
                                                 Cisco Systems, Inc.
                                                          K. Patel
                                                        Arrcus, Inc.
                                                            S. Ray
                                             Individual Contributor
                                                           J. Dong
                                              Huawei Technologies
                                                   March 24, 2019

        BGP-LS extensions for Segment Routing BGP Egress Peer Engineering
               draft-ietf-idr-bgpls-segment-routing-epe-18

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 segments allow
   steering a flow through any topological path and service chain while
   maintaining per-flow state only at the ingress node of the SR domain.

   This document describes an extension to BGP Link State (BGP-LS) for
   advertisement of BGP Peering Segments along with their BGP peering
   node information so that efficient BGP Egress Peer Engineering (EPE)
   policies and strategies can be computed based on Segment Routing.

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

Copyright Notice

Table of Contents

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 segments allows to enforce a flow through any topological
path or service function while maintaining per-flow state only at the
ingress node of the SR domain.

The SR architecture [RFC8402] defines three types of BGP Peering
Segments that may be instantiated at a BGP node:

o Peer Node Segment (PeerNode SID) : instruction to steer to a
   specific peer node

o Peer Adjacency Segment (PeerAdj SID) : instruction to steer over a
   specific local interface towards a specific peer node

o Peer Set Segment (PeerSet SID) : instruction to load-balance to a
   set of specific peer nodes

SR can be directly applied to either an MPLS dataplane (SR/MPLS) with
no change on the forwarding plane or to a modified IPv6 forwarding
plane (SRv6).

This document describes extensions to the BGP Link State NLRI (BGP-LS
NLRI) and the BGP-LS Attribute defined for BGP-LS [RFC7752] for
advertising BGP peering segments from a BGP node along with its
peering topology information (i.e. its peers, interfaces, and peering
ASs) to enable computation of efficient BGP Egress Peer Engineering
(BGP-EPE) policies and strategies using the SR/MPLS dataplane. The
corresponding extensions for SRv6 are specified in
[I-D.dawra-idr-bgpls-srv6-ext].

[I-D.ietf-spring-segment-routing-central-epe] illustrates a
centralized controller based BGP Egress Peer Engineering solution
involving SR path computation using the BGP Peering Segments. This
use-case comprises of a centralized controller that learns the BGP
Peering SIDs via BGP-LS and then uses this information to program a
BGP-EPE policy at any node in the domain to perform traffic steering
via a specific BGP egress node to a specific EBGP peer(s) optionally
also over a specific interface. The BGP-EPE policy can be realized
using the SR Policy framework
[I-D.ietf-spring-segment-routing-policy].

This document introduces a new BGP-LS Protocol-ID for BGP and defines
new BGP-LS Node and Link Descriptor TLVs to facilitate advertising

BGP-LS Link NLRI to represent the BGP peering topology. Further, it
specifies the BGP-LS Attribute TLVs for advertisement of the BGP
Peering Segments (i.e. PeerNode SID, PeerAdj SID, and PeerSet SID)
to be advertised in the same BGP-LS Link NLRI.

2. BGP Peering Segments

As described in [RFC8402], a BGP-EPE enabled Egress PE node
instantiates SR Segments corresponding to its attached peers. These
segments 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 to a specific EBGP 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.

Each BGP session MUST be described by a PeerNode SID. The
description of the BGP session MAY be augmented by additional PeerAdj
SIDs. Finally, multiple PeerNode SIDs or PeerAdj SIDs MAY be part of
the same group/set in order to group EPE resources under a common
PeerSet SID. These BGP Peering SIDs and their encoding are described
in detail in Section 4.

The following BGP Peering SIDs need to be instantiated on a BGP
router for each of its BGP peer sessions that are enabled for Egress
Peer Engineering:

o One PeerNode SID MUST be instantiated to describe the BGP peer
   session.

o One or more PeerAdj SID MAY be instantiated corresponding to the
   underlying link(s) to the directly connected BGP peer session.

o A PeerSet SID MAY be instantiated and additionally associated and
   shared between one or more PeerNode SIDs or PeerAdj SIDs.

While an egress point in a 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. However, unlike EBGP sessions which are
generally between directly connected BGP routers which are also along
the traffic forwarding path, IBGP peer sessions may be setup to BGP
routers which are not in the forwarding path. As such, when the IBGP
design includes sessions with route-reflectors, a BGP router SHOULD

NOT instantiate a BGP Peering SID for those sessions to peer nodes
which are not in the forwarding path since the purpose of BGP Peering
SID is to steer traffic to that specific peers. Thus, the
applicability for IBGP peering may be limited to only those
deployments where the IBGP peer is also along the forwarding data
path.

Any BGP Peering SIDs instantiated on the node are advertised via BGP-
LS Link NLRI type as described in the sections below. An
illustration of the BGP Peering SIDs' allocations in a reference BGP
peering topology along with the information carried in the BGP-LS
Link NLRI and its corresponding BGP-LS Attribute are described in
[I-D.ietf-spring-segment-routing-central-epe].

3. BGP-LS NLRI Advertisement for BGP Protocol

his section describes the BGP-LS NLRI encodings that describe the BGP
peering and link connectivity between BGP routers.

This document specifies the advertisement of BGP peering topology
information via BGP-LS Link NLRI type which requires use of a new
BGP-LS Protocol-ID.

```
+-------------+----------------------------------+
| Protocol-ID | NLRI information source protocol |
+-------------+----------------------------------+
|      7      | BGP                              |
+-------------+----------------------------------+
```

Table 1: BGP-LS Protocol Identifier for BGP

The use of a new Protocol-ID allows separation and differentiation
between the BGP-LS NLRIs carrying BGP information from the BGP-LS
NLRIs carrying IGP link-state information defined in [RFC7752].

The BGP Peering information along with their Peering Segments are
advertised using BGP-LS Link NLRI type with the Protocol-ID set to
BGP. The BGP-LS Link NLRI type uses the Descriptor TLVs and BGP-LS
Attribute TLVs as defined in [RFC7752]. In order to correctly
describe BGP nodes, new TLVs are defined in this section.

[RFC7752] defines Link NLRI Type is 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
+-+-+-+-+-+-+-+-+
| Protocol-ID |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          Identifier                           |
|                          (64 bits)                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
//        Local Node Descriptors                               //
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
//        Remote Node Descriptors                              //
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
//        Link Descriptors                                     //
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
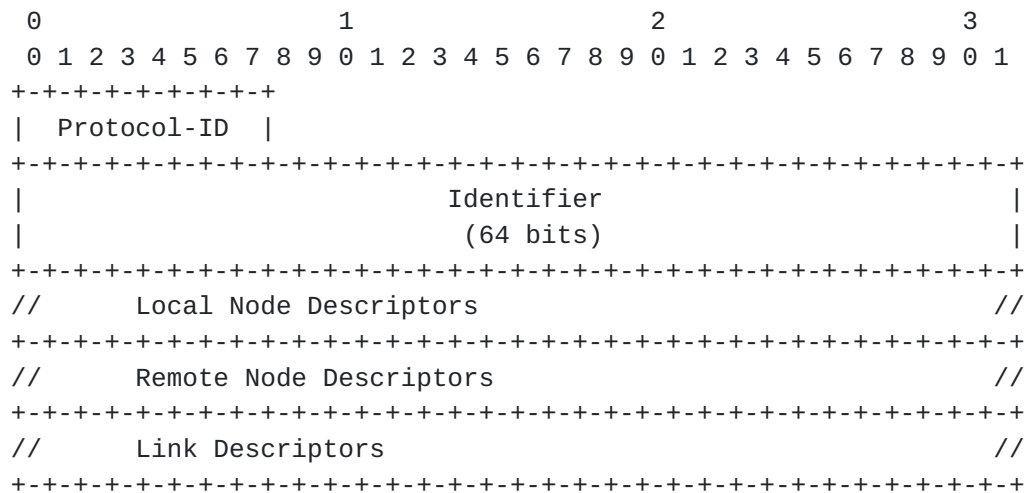
                    Figure 1: BGP-LS Link NLRI

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

3.1. BGP Router-ID and Member AS Number

   Two new Node Descriptors TLVs are defined in this document:

   o BGP Router Identifier (BGP Router-ID):

         Type: 516

         Length: 4 octets

         Value: 4 octet unsigned non-zero integer representing the BGP
         Identifier as defined in [RFC6286].

   o Member-AS Number (Member-ASN)

         Type: 517

         Length: 4 octets

         Value: 4 octet unsigned non-zero integer representing the
         Member-AS Number [RFC5065].

3.2. Mandatory BGP Node Descriptors

   The following Node Descriptors TLVs MUST be included in BGP-LS NLRI
   as Local Node Descriptors when distributing BGP information:

   o BGP Router-ID (TLV 516), which contains a valid BGP Identifier of
      the local BGP node.

o Autonomous System Number (TLV 512), which contains the ASN or AS
   Confederation Identifier (ASN) [RFC5065], if confederations are
   used, of the local BGP node.

Note that [RFC6286] (section 2.1) requires the BGP identifier
(Router-ID) to be unique within an Autonomous System and non-zero.
Therefore, the <ASN, BGP Router-ID> tuple is globally unique. Their
use in the Node Descriptor helps map Link-State NLRIs with BGP
protocol-ID to a unique BGP router in the administrative domain where
BGP-LS is enabled.

The following Node Descriptors TLVs MUST be included in BGP-LS Link
NLRI as Remote Node Descriptors when distributing BGP information:

o BGP Router-ID (TLV 516), which contains the valid BGP Identifier
   of the peer BGP node.

o Autonomous System Number (TLV 512), which contains the ASN or the
   AS Confederation Identifier (ASN) [RFC5065], if confederations are
   used, of the peer BGP node.

3.3. Optional BGP Node Descriptors

The following Node Descriptors TLVs MAY be included in BGP-LS NLRI as
Local Node Descriptors when distributing BGP information:

o Member-ASN (TLV 517), which contains the ASN of the confederation
   member (i.e. Member-AS Number), if BGP confederations are used,
   of the local BGP node.

o Node Descriptors as defined in [RFC7752].

The following Node Descriptors TLVs MAY be included in BGP-LS Link
NLRI as Remote Node Descriptors when distributing BGP information:

o Member-ASN (TLV 517), which contains the ASN of the confederation
   member (i.e. Member-AS Number), if BGP confederations are used,
   of the peer BGP node.

o Node Descriptors as defined in [RFC7752].

4. BGP-LS Attributes for BGP Peering Segments

This section defines the BGP-LS Attributes corresponding to the
following BGP Peer Segment SIDs:

   Peer Node Segment Identifier (PeerNode SID)

Peer Adjacency Segment Identifier (PeerAdj SID)

Peer Set Segment Identifier (PeerSet SID)

The following new BGP-LS Link attributes TLVs are defined for use
with BGP-LS Link NLRI for advertising BGP Peering SIDs:

```
+----------+--------------------------+----------+
| TLV Code | Description              |  Length  |
|  Point   |                          |          |
+----------+--------------------------+----------+
|    1101  | PeerNode SID             | variable |
|    1102  | PeerAdj SID              | variable |
|    1103  | PeerSet SID              | variable |
+----------+--------------------------+----------+
```

Figure 2: BGP-LS TLV code points for BGP-EPE

PeerNode SID, PeerAdj SID, and PeerSet SID have all the same format
defined here below:

```
 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         |    Weight     |            Reserved           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SID/Label/Index (variable)                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: BGP Peering SIDs TLV Format

o Type: 1101, 1102 or 1103 as listed in Figure 2.

o Length: variable. Valid values are either 7 or 8 based on the
   whether the encoding is done as a SID Index or a label.

o Flags: one octet of flags with the following definition:

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|V|L|B|P| Rsvd  |
+-+-+-+-+-+-+-+-+
```

Figure 4: Peering SID TLV Flags Format

* V-Flag: Value flag. If set, then the SID carries a label
  value. By default the flag is SET.

* L-Flag: Local Flag. If set, then the value/index carried by
  the SID has local significance. By default the flag is SET.

* B-Flag: Backup Flag. If set, the SID refers to a path that is
  eligible for protection using fast re-route (FRR). The
  computation of the backup forwarding path and its association
  with the BGP Peering SID forwarding entry is implementation
  specific. [I-D.ietf-spring-segment-routing-central-epe]
  section 3.6 discusses some of the possible ways of identifying
  backup paths for BGP Peering SIDs.

* P-Flag: Persistent Flag: If set, the SID is persistently
  allocated, i.e., the SID value remains consistent across router
  restart and session/interface flap.

* Rsvd bits: Reserved for future use and 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 [RFC8402].

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 Segment Routing
     Global Block (SRGB) [RFC8402] 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].

The values of the PeerNode SID, PeerAdj SID, and PeerSet SID Sub-TLVs
SHOULD be persistent across router restart.

When enabled for Egress Peer Engineering, the BGP router MUST include
the PeerNode SID TLV in the BGP-LS Attribute for the BGP-LS Link NLRI
corresponding to its BGP peering sessions. The PeerAdj SID and
PeerSet SID TLVs MAY be included in the BGP-LS Attribute for the BGP-
LS Link NLRI.

Additional BGP-LS Link Attribute TLVs, as defined in [RFC7752] MAY be included with the BGP-LS Link NLRI in order to advertise the characteristics of the peering link.

4.1. Advertisement of the PeerNode SID

The PeerNode SID TLV includes a SID associated with the BGP peer node that is described by a BGP-LS Link NLRI as specified in Section 3.

The PeerNode SID, at the BGP node advertising it, has the following semantics (as defined in [RFC8402]):

o SR operation: NEXT.

o Next-Hop: the connected peering node to which the segment is associated.

The PeerNode SID is advertised with a BGP-LS Link NLRI, where:

o Local Node Descriptors include:

   * Local BGP Router-ID (TLV 516) of the BGP-EPE enabled egress PE.

   * Local ASN (TLV 512).

o Remote Node Descriptors include:

   * Peer BGP Router-ID (TLV 516) (i.e.: the peer BGP ID used in the BGP session)

   * Peer ASN (TLV 512).

o Link Descriptors include the addresses used by the BGP session encoded using TLVs as defined in [RFC7752]:

   * IPv4 Interface Address (TLV 259) contains the BGP session IPv4 local address.

   * IPv4 Neighbor Address (TLV 260) contains the BGP session IPv4 peer address.

   * IPv6 Interface Address (TLV 261) contains the BGP session IPv6 local address.

   * IPv6 Neighbor Address (TLV 262) contains the BGP session IPv6 peer address.

o Link Attribute TLVs include the PeerNode SID TLV as defined in
   Figure 3.

4.2. Advertisement of the PeerAdj SID

The PeerAdj SID TLV includes a SID associated with the underlying
link to the BGP peer node that is described by a BGP-LS Link NLRI as
specified in Section 3.

The PeerAdj SID, at the BGP node advertising it, has the following
semantics (as defined in [RFC8402]):

o SR operation: NEXT.

o Next-Hop: the interface peer address.

The PeerAdj SID is advertised with a BGP-LS Link NLRI, where:

o Local Node Descriptors include:

   * Local BGP Router-ID (TLV 516) of the BGP-EPE enabled egress PE.

   * Local ASN (TLV 512).

o Remote Node Descriptors include:

   * Peer BGP Router-ID (TLV 516) (i.e. the peer BGP ID used in the
      BGP session).

   * Peer ASN (TLV 512).

o Link Descriptors MUST include the following TLV, as defined in
   [RFC7752]:

   * Link Local/Remote Identifiers (TLV 258) contains the 4-octet
      Link Local Identifier followed by the 4-octet Link Remote
      Identifier. The value 0 is used by default when the link
      remote identifier is unknown.

o Additional Link Descriptors TLVs, as defined in [RFC7752], MAY
   also be included to describe the addresses corresponding to the
   link between the BGP routers:

   * 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 TLVs include the PeerAdj SID TLV as defined in
   Figure 3.

4.3. Advertisement of the PeerSet SID

The PeerSet SID TLV includes a SID that is shared amongst BGP peer
nodes or the underlying links that are described by BGP-LS Link NLRI
as specified in Section 3.

The PeerSet SID, at the BGP node advertising it, has the following
semantics (as defined in [RFC8402]):

o SR operation: NEXT.

o Next-Hop: load balance across any connected interface to any peer
   in the associated peer set.

The PeerSet SID TLV containing the same SID value (encoded as defined
in Figure 3) is included in the BGP-LS Attribute for all of the BGP-
LS Link NLRI corresponding to the PeerNode or PeerAdj segments
associated with the peer set.

5. 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.

5.1. New BGP-LS Protocol-ID

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

```
+--------------------------------------------------------+
| Codepoint | Description |          Status              |
+--------------------------------------------------------+
|    7      | BGP         | Early Allocation by IANA      |
+--------------------------------------------------------+
```

Figure 5: BGP Protocol Codepoint

5.2. 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".

The following new Node Descriptors TLVs are defined:

```
+----------------------------------------------------------------------+
| Codepoint | Description              |          Status               |
+----------------------------------------------------------------------+
|    516    | BGP Router-ID            | Early Allocation by IANA       |
|    517    | BGP Confederation Member | Early Allocation by IANA       |
+-----------+----------------------------------------------------------+
```

Figure 6: BGP-LS Descriptor TLVs Codepoints

The following new Link Attribute TLVs are defined:

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

Figure 7: BGP-LS Attribute TLVs Codepoints

6. **Manageability Considerations**

   The new protocol extensions introduced in this document augment the
   existing IGP topology information BGP-LS distribution [RFC7752] by
   adding support for distribution of BGP peering topology information.
   As such, the Manageability Considerations section of [RFC7752]
   applies to these new extensions as well.

   Specifically, the malformed Link-State NLRI and BGP-LS Attribute
   tests for syntactic checks in the Fault Management section of
   [RFC7752] now apply to the TLVs defined in this document. The
   semantic or content checking for the TLVs specified in this document
   and their association with the BGP-LS NLRI types or their associated
   BGP-LS Attributes is left to the consumer of the BGP-LS information
   (e.g. an application or a controller) and not the BGP protocol.

   A consumer of the BGP-LS information retrieves this information from
   a BGP protocol component, that is doing the signaling over a BGP-LS
   session, via some APIs or a data model (refer Section 1 and 2 of
   [RFC7752]). The handling of semantic or content errors by the
   consumer would be dictated by the nature of its application usage and
   hence is beyond the scope of this document. It may be expected that
   an error detected in the NLRI descriptor TLVs would result in that
   specific NLRI update being unusable and hence its update to be
   discarded along with an error log. While an error in Attribute TLVs
   would result in only that specific attribute being discarded with an
   error log.

   The operator MUST be provided with the options of configuring,
   enabling, and disabling the advertisement of each of the Peer-Node-
   SID, Peer-Adj-SID, and Peer-Set-SID as well as control of 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.

   BGP Peering Segments are associated with the normal BGP routing
   peering sessions. However, the BGP peering information along with
   these Peering Segments themselves are advertised via a distinct BGP-
   LS peering session. It is expected that this isolation as described
   in [RFC7752] is followed when advertising BGP peering topology
   information via BGP-LS.

   BGP-EPE functionality enables the capability for instantiation of an
   SR path for traffic engineering a flow via an egress BGP router to a
   specific peer, bypassing the normal BGP best path routing for that
   flow and any routing policies implemented in BGP on that egress BGP
   router. As with any traffic engineering solution, the controller or
   application implementing the policy needs to ensure that there is no

looping or mis-routing of traffic. Traffic counters corresponding to
the MPLS label of the BGP Peering SID on the router would indicate
the traffic being forwarded based on the specific EPE path.
Monitoring these counters and the flows hitting the corresponding
MPLS forwarding entry would help identify issues, if any, with
traffic engineering over the EPE paths. Errors in the encoding or
decoding of the SR information in the TLVs defined in this document
may result in the unavailability of such information to a Centralized
EPE Controller or incorrect information being made available to it.
This may result in the controller not being able to perform the
desired SR based optimization functionality or to perform it in an
unexpected or inconsistent manner. The handling of such errors by
applications like such a controller may be implementation specific
and out of scope of this document.

7. Security Considerations

   [RFC7752] defines BGP-LS NLRI to which the extensions defined in this
   document apply. The Security Considerations section of [RFC7752]
   also applies to these extensions. The procedures and new TLVs
   defined in this document, by themselves, do not affect the BGP-LS
   security model discussed in [RFC7752].

   BGP-EPE enables engineering of traffic when leaving the
   administrative domain via an egress BGP router. Therefore precaution
   is necessary to ensure that the BGP peering information collected via
   BGP-LS is limited to specific consumers in a secure manner. By
   default, Segment Routing operates within a trusted domain [RFC8402]
   and its security considerations also apply to BGP Peering Segments.
   The BGP-EPE policies are expected to be used entirely within this
   trusted SR domain (e.g. between multiple AS/domains within a single
   provider network).

   The isolation of BGP-LS peering sessions is also required to ensure
   that BGP-LS topology information (including the newly added BGP
   peering topology) is not advertised to an external BGP peering
   session outside an administrative domain.

8. Contributors

   Mach (Guoyi) Chen
   Huawei Technologies
   China

   Email: mach.chen@huawei.com

      Acee Lindem
      Cisco Systems Inc.
      US

      Email: acee@cisco.com

9. Acknowledgements

   The authors would like to thank Jakob Heitz, Howard Yang, Hannes
   Gredler, Peter Psenak, Arjun Sreekantiah and Bruno Decraene for their
   feedback and comments. Susan Hares helped in improving the clarity
   of the document with her substantial contributions during her
   shepherd's review. The authors would also like to thank Alvaro
   Retana for his extensive review and comments which helped correct
   issues and improve the document.

10. References

10.1. Normative References

   [I-D.ietf-idr-bgp-ls-segment-routing-ext]
              Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H.,
              and M. Chen, "BGP Link-State extensions for Segment
              Routing", draft-ietf-idr-bgp-ls-segment-routing-ext-12
              (work in progress), March 2019.

   [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>.

   [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>.

   [RFC6286] Chen, E. and J. Yuan, "Autonomous-System-Wide Unique BGP
              Identifier for BGP-4", RFC 6286, DOI 10.17487/RFC6286,
              June 2011, <https://www.rfc-editor.org/info/rfc6286>.

   [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>.

   [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L.,
              Decraene, B., Litkowski, S., and R. Shakir, "Segment
              Routing Architecture", RFC 8402, DOI 10.17487/RFC8402,
              July 2018, <https://www.rfc-editor.org/info/rfc8402>.

10.2. Informative References

   [I-D.dawra-idr-bgpls-srv6-ext]
              Dawra, G., Filsfils, C., Talaulikar, K., Chen, M.,
              daniel.bernier@bell.ca, d., Uttaro, J., Decraene, B., and
              H. Elmalky, "BGP Link State Extensions for SRv6", draft-
              dawra-idr-bgpls-srv6-ext-05 (work in progress), March
              2019.

   [I-D.ietf-spring-segment-routing-central-epe]
              Filsfils, C., Previdi, S., Dawra, G., Aries, E., and D.
              Afanasiev, "Segment Routing Centralized BGP Egress Peer
              Engineering", draft-ietf-spring-segment-routing-central-
              epe-10 (work in progress), December 2017.

   [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-02 (work in progress), October 2018.

Authors' Addresses

   Stefano Previdi
   Individual

   Email: stefano@previdi.net

   Ketan Talaulikar (editor)
   Cisco Systems, Inc.
   India

   Email: ketant@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

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