Network Working Group   C. Filsfils, Ed.
Internet-Draft   S. Previdi
Intended status: Informational   G. Dawra, Ed.
Expires: May 1, 2018   Cisco Systems, Inc.
                                                          E. Aries
                                                   Juniper Networks
                                                      D. Afanasiev
                                                            Yandex
                                                  October 28, 2017

       Segment Routing Centralized BGP Egress Peer Engineering
           draft-ietf-spring-segment-routing-central-epe-07

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 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 a
   minor extension to the existing link-state routing protocols.

   This document illustrates the application of Segment Routing to solve
   the BGP Egress Peer Engineering (BGP-EPE) requirement. The SR-based
   BGP-EPE solution allows a centralized (Software Defined Network, SDN)
   controller to program any egress peer policy at ingress border
   routers or at hosts within the domain.

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

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 May 1, 2018.

Copyright Notice

Table of Contents

1. Introduction

The document is structured as follows:

o Section 1 states the BGP-EPE problem statement and provides the
   key references.

o Section 2 defines the different BGP Peering Segments and the
   semantic associated to them.

o Section 3 describes the automated allocation of BGP Peering SID's
   by the BGP-EPE enabled egress border router and the automated
   signaling of the external peering topology and the related BGP
   Peering SID's to the collector
   [I-D.ietf-idr-bgpls-segment-routing-epe].

o Section 4 overviews the components of a centralized BGP-EPE
   controller. The definition of the BGP-EPE controller is outside
   the scope of this document.

o Section 5 overviews the methods that could be used by the
   centralized BGP-EPE controller to implement a BGP-EPE policy at an
   ingress border router or at a source host within the domain. The
   exhaustive definition of all the means to program an BGP-EPE input
   policy is outside the scope of this document.

For editorial reasons, the solution is described with IPv6 addresses
and MPLS SIDs. This solution is equally applicable to IPv4 with MPLS
SIDs and also to IPv6 with native IPv6 SIDs.

1.1. Problem Statement

The BGP-EPE problem statement is defined in [RFC7855].

A centralized controller should be able to instruct an ingress
Provider Edge router (PE) or a content source within the domain to

use a specific egress PE and a specific external interface/neighbor
to reach a particular destination.

Let's call this solution "BGP-EPE" for "BGP Egress Peer Engineering".
The centralized controller is called the "BGP-EPE Controller". The
egress border router where the BGP-EPE traffic steering functionality
is implemented is called a BGP-EPE enabled border router. The input
policy programmed at an ingress border router or at a source host is
called a BGP-EPE policy.

The requirements that have motivated the solution described in this
document are listed here below:

o The solution MUST apply to the Internet use-case where the
   Internet routes are assumed to use IPv4 unlabeled or IPv6
   unlabeled. It is not required to place the Internet routes in a
   VRF and allocate labels on a per route, or on a per-path basis.

o The solution MUST support any deployed iBGP schemes (RRs,
   confederations or iBGP full meshes).

o The solution MUST be applicable to both routers with external and
   internal peers.

o The solution SHOULD minimize the need for new BGP capabilities at
   the ingress PEs.

o The solution MUST accommodate an ingress BGP-EPE policy at an
   ingress PE or directly at a source host within the domain.

o The solution MAY support automated Fast Reroute (FRR) and fast
   convergence mechanisms.

The following reference diagram is used throughout this document.

```
+---------+       +------+
|         |       |      |
|         |       |      |
|    H    B------D       G
|         | +---/| AS 2 |\   +------+
|         |/     +------+ \ |      |---L/8
A    AS1   C---+            \|       |
|         |\\   \  +------+ /| AS 4 |---M/8
|         | \\   +-E       |/ +------+
|    X    |  \\    |        K
|         |    +===F AS 3 |
+---------+       +------+
```
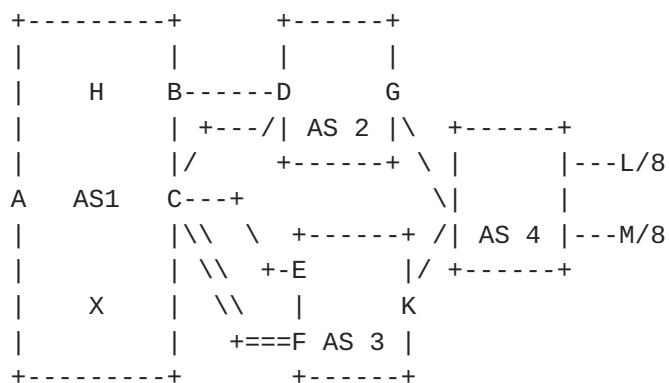
Figure 1: Reference Diagram

IP addressing:

o C's interface to D: 2001:db8:cd::c/64, D's interface:
   2001:db8:cd::d/64

o C's interface to E: 2001:db8:ce::c/64, E's interface:
   2001:db8:ce::e/64

o C's upper interface to F: 2001:db8:cf1::c/64, F's interface:
   2001:db8:cf1::f/64

o C's lower interface to F: 2001:db8:cf2::c/64, F's interface:
   2001:db8:cf2::f/64

o BGP router-ID of C: 192.0.2.3

o BGP router-ID of D: 192.0.2.4

o BGP router-ID of E: 192.0.2.5

o BGP router-ID of F: 192.0.2.6

o Loopback of F used for eBGP multi-hop peering to C:
   2001:db8:f::f/128

o C's loopback is 2001:db8:c::c/128 with SID 64

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:ce::e (E)

o Multi-hop eBGP peering with F on IP address 2001:db8:f::f (F)

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

o Static route to 2001:db8:f::f/128 via 2001:db8:cf1::f

o Static route to 2001:db8:f::f/128 via 2001:db8:cf2::f

C is configured with local policy that defines a BGP PeerSet as the
set of peers (2001:db8:ce::e for E and 2001:db8:f::f for F)

X is the BGP-EPE controller within AS1 domain.

H is a content source within AS1 domain.

2. **BGP Peering Segments**

   As defined in [I-D.ietf-spring-segment-routing], certain segments are
   defined by BGP-EPE capable node and corresponding to its attached
   peers. These segments are called BGP peering segments or BGP Peering
   SIDs. They enable the expression of source-routed inter-domain
   paths.

   An ingress border router of an AS may compose a list of segments 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 Egress Peering Engineering policy applied at an
   ingress EPE involves two segments: the Node SID of the chosen egress
   EPE and then the BGP Peering Segment for the chosen egress EPE peer
   or peering interface.

   [I-D.ietf-spring-segment-routing] defines three types of BGP peering
   segments/SIDs: PeerNode SID, PeerAdj SID and PeerSet SID.

   A Peer Node Segment is a segment describing a peer, including the SID
   (PeerNode SID) allocated to it.

   A Peer Adjacency Segment is a segment describing a link, including
   the SID (PeerAdj SID) allocated to it.

   A Peer Set Segment is a segment describing a link or a node that is
   part of the set, including the SID (PeerSet SID) allocated to the
   set.

3. **Distribution of Topology and TE Information using BGP-LS**

   In ships-in-the-night mode with respect to the pre-existing iBGP
   design, a BGP-LS [RFC7752] session is established between the BGP-EPE
   enabled border router and the BGP-EPE controller.

   As a result of its local configuration and according to the behavior
   described in [I-D.ietf-idr-bgpls-segment-routing-epe], node C
   allocates the following BGP Peering Segments
   ([I-D.ietf-spring-segment-routing]):

   o A PeerNode segment for each of its defined peer (D: 1012, E: 1022
      and F: 1052).

   o A PeerAdj segment for each recursing interface to a multi-hop peer
      (e.g.: the upper and lower interfaces from C to F in figure 1).

o A PeerSet segment to the set of peers (E and F). In this case the
   PeerSet represents a set of peers (E, F) belonging to the same AS
   (AS 3).

C programs its forwarding table accordingly:

```
Incoming               Outgoing
Label      Operation   Interface
------------------------------------
1012         POP     link to D
1022         POP     link to E
1032         POP     upper link to F
1042         POP     lower link to F
1052         POP     load balance on any link to F
1060         POP     load balance on any link to E or to F
```

C signals the related BGP-LS NLRI's to the BGP-EPE controller. Each
such BGP-LS route is described in the following subsections according
to the encoding details defined in
[I-D.ietf-idr-bgpls-segment-routing-epe].

3.1. PeerNode SID to D

Descriptors:

o Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier):
   192.0.2.3, AS1, 1000

o Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.4, AS2

o Link Descriptors (IPv6 Interface Address, IPv6 Neighbor Address):
   2001:db8:cd::c, 2001:db8:cd::d

Attributes:

o PeerNode SID: 1012

3.2. PeerNode SID to E

Descriptors:

o Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier)):
   192.0.2.3, AS1, 1000

o Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.5, AS3

o Link Descriptors (IP interface address, neighbor IP address):
   2001:db8:ce::c, 2001:db8:ce::e

Attributes:

o PeerNode SID: 1022

o PeerSetSID: 1060

o Link Attributes: see section 3.3.2 of [RFC7752]

3.3. PeerNode SID to F

Descriptors:

o Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier)):
   192.0.2.3, AS1, 1000

o Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.6, AS3

o Link Descriptors (IP interface address, neighbor IP address):
   2001:db8:c::c, 2001:db8:f::f

Attributes:

o PeerNode SID: 1052

o PeerSetSID: 1060

3.4. First PeerAdj to F

Descriptors:

o Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier)):
   192.0.2.3, AS1, 1000

o Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.6, AS3

o Link Descriptors (IP interface address, neighbor IP address):
   2001:db8:cf1::c, 2001:db8:cf1::f

Attributes:

o PeerAdj-SID: 1032

o LinkAttributes: see section 3.3.2 of [RFC7752]

3.5. Second PeerAdj to F

Descriptors:

o Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier)):
   192.0.2.3 , AS1

o Remote Node Descriptors (peer router-ID, peer ASN): 192.0.2.6, AS3

o Link Descriptors (IP interface address, neighbor IP address):
   2001:db8:cf2::c, 2001:db8:cf2::f

Attributes:

o PeerAdj-SID: 1042

o LinkAttributes: see [section 3.3.2 of [RFC7752]](#)

3.6. Fast Reroute (FRR)

A BGP-EPE enabled border router MAY allocate a FRR backup entry on a
per BGP Peering SID basis. One example is as follows:

o PeerNode SID

   1. If multi-hop, backup via the remaining PeerADJ SIDs (if
       available) to the same peer.

   2. Else backup via another PeerNode SID to the same AS.

   3. Else pop the PeerNode SID and perform an IP lookup.

o PeerAdj SID

   1. If to a multi-hop peer, backup via the remaining PeerADJ SIDs
       (if available) to the same peer.

   2. Else backup via a PeerNode SID to the same AS.

   3. Else pop the PeerNode SID and perform an IP lookup.

o PeerSet SID

   1. Backup via remaining PeerNode SIDs in the same PeerSet.

   2. Else pop the PeerNode SID and IP lookup.

Let's illustrate different types of possible backups using the
reference diagram and considering the Peering SIDs allocated by C.

PeerNode SID 1052, allocated by C for peer F:

o Upon the failure of the upper connected link CF, C can reroute all
   the traffic onto the lower CF link to the same peer (F).

PeerNode SID 1022, allocated by C for peer E:

o Upon the failure of the connected link CE, C can reroute all the
   traffic onto the link to PeerNode SID 1052 (F).

PeerNode SID 1012, allocated by C for peer D:

o Upon the failure of the connected link CD, C can pop the PeerNode
   SID and lookup the IP destination address in its FIB and route
   accordingly.

PeerSet SID 1060, allocated by C for the set of peers E and F:

o Upon the failure of a connected link in the group, the traffic to
   PeerSet SID 1060 is rerouted on any other member of the group.

For specific business reasons, the operator might not want the
default FRR behavior applied to a PeerNode SID or any of its
dependent PeerADJ SID.

The operator should be able to associate a specific backup PeerNode
SID for a PeerNode SID: e.g., 1022 (E) must be backed up by 1012 (D)
which overrules the default behavior which would have preferred F as
a backup for E.

4. BGP-EPE Controller

In this section, Let's provide a non-exhaustive set of inputs that a
BGP-EPE controller would likely collect such as to perform the BGP-
EPE policy decision.

The exhaustive definition is outside the scope of this document.

4.1. Valid Paths From Peers

The BGP-EPE controller should collect all the BGP paths (i.e.: IP
destination prefixes) advertised by all the BGP-EPE enabled border
router.

This could be realized by setting an iBGP session with the BGP-EPE enabled border router, with the router configured to advertise all paths using BGP add-path [RFC7911] and the original next-hop preserved.

In this case, C would advertise the following Internet routes to the BGP-EPE controller:

o NLRI <2001:db8:abcd::/48>, next-hop 2001:db8:cd::d, AS Path {AS 2, 4}

   * X (i.e.: the BGP-EPE controller) knows that C receives a path to 2001:db8:abcd::/48 via neighbor 2001:db8:cd::d of AS2.

o NLRI <2001:db8:abcd::/48>, next-hop 2001:db8:ce::e, AS Path {AS 3, 4}

   * X knows that C receives a path to 2001:db8:abcd::/48 via neighbor 2001:db8:ce::e of AS2.

o NLRI <2001:db8:abcd::/48>, next-hop 2001:db8:f::f, AS Path {AS 3, 4}

   * X knows that C has an eBGP path to 2001:db8:abcd::/48 via AS3 via neighbor 2001:db8:f::f

An alternative option would be for a BGP-EPE collector to use BGP Monitoring Protocol (BMP) [RFC7854] to track the Adj-RIB-In of BGP-EPE enabled border routers.

4.2. Intra-Domain Topology

The BGP-EPE controller should collect the internal topology and the related IGP SIDs.

This could be realized by collecting the IGP LSDB of each area or running a BGP-LS session with a node in each IGP area.

4.3. External Topology

Thanks to the collected BGP-LS routes described in section 2, the BGP-EPE controller is able to maintain an accurate description of the egress topology of node C. Furthermore, the BGP-EPE controller is able to associate BGP Peering SIDs to the various components of the external topology.

4.4. SLA characteristics of each peer

   The BGP-EPE controller might collect SLA characteristics across
   peers. This requires an BGP-EPE solution as the SLA probes need to
   be steered via non-best-path peers.

   Unidirectional SLA monitoring of the desired path is likely required.
   This might be possible when the application is controlled at the
   source and the receiver side. Unidirectional monitoring dissociates
   the SLA characteristic of the return path (which cannot usually be
   controlled) from the forward path (the one of interest for pushing
   content from a source to a consumer and the one which can be
   controlled).

   Alternatively, Extended Metrics, as defined in [RFC7810] could also
   be advertised using BGP-LS ([I-D.ietf-idr-te-pm-bgp]).

4.5. Traffic Matrix

   The BGP-EPE controller might collect the traffic matrix to its peers
   or the final destinations. IPFIX [RFC7011] is a likely option.

   An alternative option consists in collecting the link utilization
   statistics of each of the internal and external links, also available
   in the current definition of [RFC7752].

4.6. Business Policies

   The BGP-EPE controller should be configured or collect (through any
   mean) business policies. The mechanisms through which these policies
   are configured or collected is outside the scope of this document.

4.7. BGP-EPE Policy

   On the basis of all these inputs (and likely others), the BGP-EPE
   Controller decides to steer some demands away from their best BGP
   path.

   The BGP-EPE policy is likely expressed as a two-entry segment list
   where the first element is the IGP prefix SID of the selected egress
   border router and the second element is a BGP Peering SID at the
   selected egress border router.

   A few examples are provided hereafter:

   o Prefer egress PE C and peer AS AS2: {64, 1012}. "64" being the SID
      of PE C as defined in Section 1.1.

o Prefer egress PE C and peer AS AS3 via eBGP peer 2001:db8:ce::e,
   {64, 1022}.

o Prefer egress PE C and peer AS AS3 via eBGP peer 2001:db8:f::f,
   {64, 1052}.

o Prefer egress PE C and peer AS AS3 via interface 2001:db8:cf2::f
   of multi-hop eBGP peer 2001:db8:f::f, {64, 1042}.

o Prefer egress PE C and any interface to any peer in the group
   1060: {64, 1060}.

Note that the first SID could be replaced by a list of segments.
This is useful when an explicit path within the domain is required
for traffic engineering purposes. For example, if the Prefix SID of
node B is 60 and the BGP-EPE controller would like to steer the
traffic from A to C via B then through the external link to peer D
then the segment list would be {60, 64, 1012}.

5. Programming an input policy

The detailed/exhaustive description of all the means to implement an
BGP-EPE policy are outside the scope of this document. A few
examples are provided in this section.

5.1. At a Host

A static IP/MPLS route can be programmed at the host H. The static
route would define a destination prefix, a next-hop and a label stack
to push. Assuming a global SRGB, at least on all access routers
connecting the hosts, the same policy can be programmed across all
hosts, which is convenient.

5.2. At a router - SR Traffic Engineering tunnel

The BGP-EPE controller can configure the ingress border router with
an SR traffic engineering tunnel T1 and a steering-policy S1 which
causes a certain class of traffic to be mapped on the tunnel T1.

The tunnel T1 would be configured to push the required segment list.

The tunnel and the steering policy could be configured via multiple
means. A few examples are given below:

o PCEP according to [I-D.ietf-pce-segment-routing] and
   [I-D.ietf-pce-pce-initiated-lsp].

o Netconf ([RFC6241]).

o Other static or ephemeral APIs

Example: at router A (Figure 1).

                     Tunnel T1: push {64, 1042}
                     IP route L/8 set next-hop T1

5.3. At a Router - RFC3107 policy route

The BGP-EPE Controller could build a RFC3107
([I-D.ietf-mpls-rfc3107bis]) route (from scratch) and send it to the
ingress router:

o NLRI: the destination prefix to engineer: e.g., L/8.

o Next-Hop: the selected egress border router: C.

o Label: the selected egress peer: 1042.

o AS path: reflecting the selected valid AS path.

o Some BGP policy to ensure it will be selected as best by the
   ingress router.

This RFC3107 policy route "overwrites" an equivalent or less-specific
"best path". As the best-path is changed, this BGP-EPE input policy
option may influence the path propagated to the upstream peer/
customers. Indeed, implementations treating the SAFI-1 and SAFI-4
routes for a given prefix as comparable would trigger a BGP WITHDRAW
of the SAFI-1 route to their BGP upstream peers.

5.4. At a Router - VPN policy route

The BGP-EPE Controller could build a VPNv4 route (from scratch) and
send it to the ingress router:

o NLRI: the destination prefix to engineer: e.g., L/8.

o Next-Hop: the selected egress border router: C.

o Label: the selected egress peer: 1042.

o Route-Target: selecting the appropriate VRF at the ingress router.

o AS path: reflecting the selected valid AS path.

o Some BGP policy to ensure it will be selected as best by the
   ingress router in the related VRF.

The related VRF must be preconfigured. A VRF fallback to the main
FIB might be beneficial to avoid replicating all the "normal"
Internet paths in each VRF.

6. IPv6 Dataplane

The described solution is applicable to IPv6, either with MPLS-based
or IPv6-Native segments. In both cases, the same three steps of the
solution are applicable:

o BGP-LS-based signaling of the external topology and BGP Peering
   Segments to the BGP-EPE controller.

o Collection of various inputs by the BGP-EPE controller to come up
   with a policy decision.

o Programming at an ingress router or source host of the desired
   BGP-EPE policy which consists in a list of segments to push on a
   defined traffic class.

7. Benefits

The BGP-EPE solutions described in this document have the following
benefits:

o No assumption on the iBGP design within AS1.

o Next-Hop-Self on the Internet routes propagated to the ingress
   border routers is possible. This is a common design rule to
   minimize the number of IGP routes and to avoid importing external
   churn into the internal routing domain.

o Consistent support for traffic engineering within the domain and
   at the external edge of the domain.

o Support both host and ingress border router BGP-EPE policy
   programming.

o BGP-EPE functionality is only required on the BGP-EPE enabled
   egress border router and the BGP-EPE controller: an ingress policy
   can be programmed at the ingress border router without any new
   functionality.

o Ability to deploy the same input policy across hosts connected to
   different routers (assuming the global property of IGP prefix
   SIDs).

8. IANA Considerations

This document does not request any IANA allocations.

9. Manageability Considerations

The BGP-EPE use-case described in this document requires BGP-LS
([RFC7752]) extensions that are described in
[I-D.ietf-idr-bgpls-segment-routing-epe]. The required extensions
consists of additional BGP-LS descriptors and TLVs that will follow
the same. Manageability functions of BGP-LS, described in [RFC7752]
also apply to the extensions required by the EPE use-case.

The operator MUST be capable of configuring, enabling, disabling the
advertisement of the EPE information 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.

10. Security Considerations

[RFC7752] defines BGP-LS NLRIs and their associated security aspects.

[I-D.ietf-idr-bgpls-segment-routing-epe] defines the BGP-LS
extensions required by the BGP-EPE mechanisms described in this
document. BGP-EPE BGP-LS extensions also include the related
security.

11. Contributors

Daniel Ginsburg substantially contributed to the content of this
document.

12. Acknowledgements

The authors would like to thank Acee Lindem for his comments and
contribution.

13. References

13.1. Normative References

[I-D.ietf-idr-bgpls-segment-routing-epe]
          Previdi, S., Filsfils, C., Patel, K., Ray, S., and J.
          Dong, "BGP-LS extensions for Segment Routing BGP Egress
          Peer Engineering", draft-ietf-idr-bgpls-segment-routing-
          epe-13 (work in progress), June 2017.

   [I-D.ietf-spring-segment-routing]
            Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B.,
            Litkowski, S., and R. Shakir, "Segment Routing
            Architecture", draft-ietf-spring-segment-routing-13 (work
            in progress), October 2017.

   [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>.

   [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>.

13.2. Informative References

   [I-D.ietf-idr-te-pm-bgp]
            Ginsberg, L., Previdi, S., Wu, Q., Gredler, H., Ray, S.,
            Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of
            IGP Traffic Engineering Performance Metric Extensions",
            draft-ietf-idr-te-pm-bgp-08 (work in progress), August
            2017.

   [I-D.ietf-mpls-rfc3107bis]
            Rosen, E., "Using BGP to Bind MPLS Labels to Address
            Prefixes", draft-ietf-mpls-rfc3107bis-04 (work in
            progress), August 2017.

   [I-D.ietf-pce-pce-initiated-lsp]
            Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP
            Extensions for PCE-initiated LSP Setup in a Stateful PCE
            Model", draft-ietf-pce-pce-initiated-lsp-11 (work in
            progress), October 2017.

   [I-D.ietf-pce-segment-routing]
            Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W.,
            and J. Hardwick, "PCEP Extensions for Segment Routing",
            draft-ietf-pce-segment-routing-10 (work in progress),
            October 2017.

   [RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed.,
            and A. Bierman, Ed., "Network Configuration Protocol
            (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011,
            <https://www.rfc-editor.org/info/rfc6241>.

   [RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken,
              "Specification of the IP Flow Information Export (IPFIX)
              Protocol for the Exchange of Flow Information", STD 77,
              RFC 7011, DOI 10.17487/RFC7011, September 2013,
              <https://www.rfc-editor.org/info/rfc7011>.

   [RFC7810] Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and
              Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions",
              RFC 7810, DOI 10.17487/RFC7810, May 2016,
              <https://www.rfc-editor.org/info/rfc7810>.

   [RFC7854] Scudder, J., Ed., Fernando, R., and S. Stuart, "BGP
              Monitoring Protocol (BMP)", RFC 7854,
              DOI 10.17487/RFC7854, June 2016,
              <https://www.rfc-editor.org/info/rfc7854>.

   [RFC7855] Previdi, S., Ed., Filsfils, C., Ed., Decraene, B.,
              Litkowski, S., Horneffer, M., and R. Shakir, "Source
              Packet Routing in Networking (SPRING) Problem Statement
              and Requirements", RFC 7855, DOI 10.17487/RFC7855, May
              2016, <https://www.rfc-editor.org/info/rfc7855>.

   [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder,
              "Advertisement of Multiple Paths in BGP", RFC 7911,
              DOI 10.17487/RFC7911, July 2016,
              <https://www.rfc-editor.org/info/rfc7911>.

Authors' Addresses

   Clarence Filsfils (editor)
   Cisco Systems, Inc.
   Brussels
   BE

   Email: cfilsfil@cisco.com

   Stefano Previdi
   Cisco Systems, Inc.
   Italy

   Email: stefano@previdi.net

Gaurav Dawra (editor)
Cisco Systems, Inc.
USA

Email: gdawra@cisco.com

Ebben Aries
Juniper Networks
1133 Innovation Way
Sunnyvale CA 94089
US

Email: exa@juniper.net

Dmitry Afanasiev
Yandex
RU

Email: fl0w@yandex-team.ru