

Internet Engineering Task Force (IETF)
Request for Comments: 9087
Category: Informational
ISSN: 2070-1721

C. Filsfils, Ed.
S. Previdi
G. Dawra, Ed.
Cisco Systems, Inc.
E. Aries
Juniper Networks
D. Afanasiev
Yandex
August 2021

Segment Routing Centralized BGP Egress Peer Engineering

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 for the enforcement of 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 data plane 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 Networking, or SDN) controller to program any egress peer policy at ingress border routers or at hosts within the domain.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9087>.

Copyright Notice

Copyright (c) 2021 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 (<https://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
1.1.	Problem Statement
1.2.	Requirements Language
2.	BGP Peering Segments
3.	Distribution of Topology and TE Information Using BGP-LS
3.1.	PeerNode SID to D
3.2.	PeerNode SID to E
3.3.	PeerNode SID to F
3.4.	First PeerAdj to F
3.5.	Second PeerAdj to F
3.6.	Fast Reroute (FRR)
4.	BGP-EPE Controller
4.1.	Valid Paths from Peers
4.2.	Intra-Domain Topology
4.3.	External Topology
4.4.	SLA Characteristics of Each Peer
4.5.	Traffic Matrix
4.6.	Business Policies
4.7.	BGP-EPE Policy
5.	Programming an Input Policy
5.1.	At a Host
5.2.	At a Router - SR Traffic-Engineering Tunnel
5.3.	At a Router - Unicast Route Labeled Using BGP (RFC 8277)
5.4.	At a Router - VPN Policy Route
6.	IPv6 Data Plane
7.	Benefits
8.	IANA Considerations
9.	Manageability Considerations
10.	Security Considerations
11.	References
11.1.	Normative References
11.2.	Informative References
	Acknowledgements
	Contributors
	Authors' Addresses

1. Introduction

The document is structured as follows:

- * Section 1 states the BGP-EPE problem statement and provides the key references.
- * Section 2 defines the different BGP Peering Segments and the semantic associated to them.
- * Section 3 describes the automated allocation of BGP Peering Segment-IDs (SIDs) by the BGP-EPE-enabled egress border router and the automated signaling of the external peering topology and the related BGP Peering SIDs to the collector [RFC9086].
- * 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.
- * 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 a 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 (PE) router 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:

- * 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 VPN Routing and Forwarding (VRF) instance and allocate labels on a per-route or per-path basis.
- * The solution MUST support any deployed Internal BGP (iBGP) schemes (Route Reflectors (RRs), confederations, or iBGP full meshes).
- * The solution MUST be applicable to both routers with external and internal peers.
- * The solution should minimize the need for new BGP capabilities at the ingress PEs.
- * The solution MUST accommodate an ingress BGP-EPE policy at an ingress PE or directly at a source within the domain.
- * The solution MAY support automated Fast Reroute (FRR) and fast convergence mechanisms.

The following reference diagram is used throughout this document.




Figure 1: Reference Diagram

IP addressing:

- * C's interface to D: 2001:db8:cd::c/64, D's interface: 2001:db8:cd::d/64
- * C's interface to E: 2001:db8:ce::c/64, E's interface: 2001:db8:ce::e/64

- * C's upper interface to F: 2001:db8:cf1::c/64, F's interface: 2001:db8:cf1::f/64
- * C's lower interface to F: 2001:db8:cf2::c/64, F's interface: 2001:db8:cf2::f/64
- * BGP router-ID of C: 192.0.2.3
- * BGP router-ID of D: 192.0.2.4
- * BGP router-ID of E: 192.0.2.5
- * BGP router-ID of F: 192.0.2.6
- * Loopback of F used for External BGP (eBGP) multi-hop peering to C: 2001:db8:f::f/128
- * C's loopback is 2001:db8:c::c/128 with SID 64

C's BGP peering:

- * Single-hop eBGP peering with neighbor 2001:db8:cd::d (D)
- * Single-hop eBGP peering with neighbor 2001:db8:ce::e (E)
- * 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:

- * Static route to 2001:db8:f::f/128 via 2001:db8:cf1::f
- * Static route to 2001:db8:f::f/128 via 2001:db8:cf2::f

C is configured with a 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 the AS1 domain.

H is a content source within the AS1 domain.

1.2. 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. BGP Peering Segments

As defined in [RFC8402], certain segments are defined by a BGP-EPE-capable node and correspond to their 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 Peer 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.

[RFC8402] defines three types of BGP Peering Segments/SIDs: PeerNode SID, PeerAdj SID, and PeerSet SID.

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

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

Peer Set Segment: 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 Border Gateway Protocol - Link State (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 [RFC9086], Node C allocates the following BGP Peering Segments [RFC8402]:

- * A PeerNode segment for each of its defined peers (D: 1012, E: 1022 and F: 1052).
- * 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).
- * 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 Label	Operation	Outgoing 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

Table 1

C signals each related BGP-LS instance of Network Layer Reachability Information (NLRI) to the BGP-EPE controller. Each such BGP-LS route is described in the following subsections according to the encoding details defined in [RFC9086].

3.1. PeerNode SID to D

Descriptors:

- * Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier): 192.0.2.3, AS1, 1000
- * Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.4, AS2

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

Attributes:

- * PeerNode SID: 1012

3.2. PeerNode SID to E

Descriptors:

- * Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier):
192.0.2.3, AS1, 1000
- * Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.5, AS3
- * Link Descriptors (IPv6 Interface Address, IPv6 Neighbor Address):
2001:db8:ce::c, 2001:db8:ce::e

Attributes:

- * PeerNode SID: 1022
- * PeerSetSID: 1060
- * Link Attributes: see Section 3.3.2 of [RFC7752]

3.3. PeerNode SID to F

Descriptors:

- * Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier):
192.0.2.3, AS1, 1000
- * Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.6, AS3
- * Link Descriptors (IPv6 Interface Address, IPv6 Neighbor Address):
2001:db8:c::c, 2001:db8:f::f

Attributes:

- * PeerNode SID: 1052
- * PeerSetSID: 1060

3.4. First PeerAdj to F

Descriptors:

- * Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier):
192.0.2.3, AS1, 1000
- * Remote Node Descriptors (BGP router-ID, ASN): 192.0.2.6, AS3
- * Link Descriptors (IPv6 Interface Address, IPv6 Neighbor Address):
2001:db8:cf1::c, 2001:db8:cf1::f

Attributes:

- * PeerAdj-SID: 1032
- * Link Attributes: see Section 3.3.2 of [RFC7752]

3.5. Second PeerAdj to F

Descriptors:

- * Local Node Descriptors (BGP router-ID, ASN, BGP-LS Identifier):
192.0.2.3 , AS1, 1000
- * Remote Node Descriptors (peer router-ID, peer ASN): 192.0.2.6, AS3
- * Link Descriptors (IPv6 Interface Address, IPv6 Neighbor Address):
2001:db8:cf2::c, 2001:db8:cf2::f

Attributes:

- * PeerAdj-SID: 1042
- * Link Attributes: see Section 3.3.2 of [RFC7752]

3.6. Fast Reroute (FRR)

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

- * PeerNode SID
 1. If multi-hop, back up via the remaining PeerADJ SIDs (if available) to the same peer.
 2. Else, back up via another PeerNode SID to the same AS.
 3. Else, pop the PeerNode SID and perform an IP lookup.
- * PeerAdj SID
 1. If to a multi-hop peer, back up via the remaining PeerADJ SIDs (if available) to the same peer.
 2. Else, back up via a PeerNode SID to the same AS.
 3. Else, pop the PeerNode SID and perform an IP lookup.
- * PeerSet SID
 1. Back up 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:

- * 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:

- * 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:

- * Upon the failure of the connected link CD, C can pop the PeerNode SID and look up the IP destination address in its FIB and route accordingly.

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

- * 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 SIDs.

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 that 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 routers.

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:

- * 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.
- * 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.
- * 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 the 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 Link-State Database (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 3, 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 Service Level Agreement (SLA) characteristics across peers. This requires a 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 that can be controlled).

Alternatively, Metric Extensions, as defined in [RFC8570], could also be advertised using BGP-LS [RFC8571].

4.5. Traffic Matrix

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

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

4.6. Business Policies

The BGP-EPE controller should be configured or collect business policies through any desired mechanisms. These mechanisms by which these policies are configured or collected are 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:

- * Prefer egress PE C and peer AS AS2: {64, 1012}. "64" being the SID of PE C as defined in Section 1.1.
- * Prefer egress PE C and peer AS AS3 via eBGP peer 2001:db8:ce::e, {64, 1022}.
- * Prefer egress PE C and peer AS AS3 via eBGP peer 2001:db8:f::f, {64, 1052}.
- * 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}.
- * 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 a 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 the same Segment Routing Global Block (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:

- * The Path Computation Element Communication Protocol (PCEP) according to [RFC8664] and [RFC8281]
- * NETCONF [RFC6241]
- * 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 - Unicast Route Labeled Using BGP (RFC 8277)

The BGP-EPE controller could build a unicast route labeled using BGP [RFC8277] (from scratch) and send it to the ingress router.

Such a route would require the following:

NLRI
the destination prefix to engineer (e.g., L/8)

Next Hop
the selected egress border router: C

Label
the selected egress peer: 1042

Autonomous System (AS) path
the selected valid AS path

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

router. Note that as discussed in Section 5 of [RFC8277], the comparison of a labeled and unlabeled unicast BGP route is implementation dependent and hence may require an implementation-specific policy on each ingress router.

This unicast route labeled using BGP [RFC8277] "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.

Such a route would require the following:

NLRI

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

Next Hop

the selected egress border router: C

Label

the selected egress peer: 1042

Route-Target

the selected appropriate VRF instance at the ingress router

AS path

the selected valid AS path

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

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

6. IPv6 Data Plane

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:

- * BGP-LS-based signaling of the external topology and BGP Peering Segments to the BGP-EPE controller.
- * Collecting, by the BGP-EPE controller, various inputs to come up with a policy decision.
- * Programming at an ingress router or source host of the desired BGP-EPE policy, which consists of 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:

- * No assumption on the iBGP design within AS1.
- * 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.

- * Consistent support for traffic engineering within the domain and at the external edge of the domain.
- * Support for both host and ingress border router BGP-EPE policy programming.
- * 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.
- * 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 has no IANA actions.

9. Manageability Considerations

The BGP-EPE use case described in this document requires BGP-LS [RFC7752] extensions that are described in [RFC9086] and that consists of additional BGP-LS descriptors and TLVs. Manageability functions of BGP-LS, described in [RFC7752], also apply to the extensions required by the EPE use case.

Additional manageability considerations are described in [RFC9086].

10. Security Considerations

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

[RFC9086] 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. References

11.1. 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>>.
- [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>>.

- [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>>.

11.2. Informative References

- [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>>.
- [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>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.

Acknowledgements

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

Contributors

Daniel Ginsburg substantially contributed to the content of this document.

Authors' Addresses

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

Email: cfilsfil@cisco.com

Stefano Previdi
Cisco Systems, Inc.
Italy

Email: stefano@previdi.net

Gaurav Dawra (editor)
Cisco Systems, Inc.
United States of America

Email: gdawra.ietf@gmail.com

Ebben Aries
Juniper Networks
1133 Innovation Way
Sunnyvale, CA 94089
United States of America

Email: exa@juniper.net

Dmitry Afanasiev
Yandex
Russian Federation

Email: fl0w@yandex-team.ru