

SPRING Working Group
Internet-Draft
Intended status: Informational
Expires: March 4, 2021

J. Dong
Huawei Technologies
S. Bryant
Futurewei Technologies
T. Miyasaka
KDDI Corporation
Y. Zhu
China Telecom
F. Qin
Z. Li
China Mobile
F. Clad
Cisco Systems
August 31, 2020

Segment Routing based Virtual Transport Network for Enhanced VPN
draft-dong-spring-sr-for-enhanced-vpn-10

Abstract

I-D.ietf-spring-resource-aware-segments describes the mechanism to associate network resource attributes to Segment Routing Identifiers (SIDs). The resource-aware SIDs can be used to build SR paths with a set of reserved network resources. In addition, the resource-aware SIDs can be used to build SR based virtual networks, which can be used as virtual underlay networks with the network topology and resource attributes required by different customers or services. Such virtual networks are called virtual transport networks (VTNs). This document describes the mechanism of using resource-aware SIDs to build SR based VTNs.

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 <https://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 March 4, 2021.

Copyright Notice

Copyright (c) 2020 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	2
2. Resource-Aware SIDs for VTN	3
2.1. SR-MPLS	3
2.2. SRv6	4
3. Procedures	5
3.1. VTN Topology and Resource Computation	5
3.2. VTN Network Resource and SID Allocation	6
3.3. Construction of SR based VTNs	8
3.4. Mapping Service to SR based VTN	9
3.5. VTN Visibility to Customer	9
4. Benefits of the Proposed Mechanism	10
5. Service Assurance	11
6. IANA Considerations	11
7. Security Considerations	11
8. Contributors	12
9. Acknowledgements	12
10. References	12
10.1. Normative References	12
10.2. Informative References	12
Authors' Addresses	15

1. Introduction

Segment Routing (SR) [[RFC8402](#)] specifies a mechanism to steer packets through an ordered list of segments. A segment is referred to by its Segment Identifier (SID). With SR, explicit source routing can be achieved without introducing per-path state into the network. While compared with RSVP-TE [[RFC3209](#)], currently SR does not have the capability of reserving network resources or identifying a set of network resources reserved for services or customers.

[[I-D.ietf-spring-resource-aware-segments](#)] extends the SR paradigm by

associating SIDs with network resource attributes. On a network segment, multiple resource-aware SIDs can be allocated, each of which represents a subset of network resources allocated to meet the requirement of some customers or services. The resource-aware SIDs can be used to build SR paths with a set of reserved network resources. In addition, the resource-aware SIDs can also be used to build SR based virtual networks with the required network topology and resource attributes. A group of resource-aware SIDs together can be used to specify the customized topology of a virtual network, and can further be used to steer the service traffic to be processed with the set of network resources allocated to the virtual network. Such virtual networks are called virtual transport networks (VTNs), which can be used as the underlay of enhanced VPN services as described in [I-D.ietf-teas-enhanced-vpn].

This document describes the mechanism of using resource-aware SIDs to build SR based VTNs. Although the procedure is illustrated using SR-MPLS, the proposed mechanism is applicable to both segment routing over MPLS data plane (SR-MPLS) and segment routing over IPv6 data plane (SRv6).

2. Resource-Aware SIDs for VTN

When SR is used as the data plane to provide multiple VTNs in one network, it is necessary that SR paths in a VTN are computed with the topology constraints, and are instantiated with the set of network resources allocated to the VTN. With the mechanism defined in [I-D.ietf-spring-resource-aware-segments], multiple SR SIDs can be allocated for each network segment, and each SID can be used to identify both the network topology and the set of network resources allocated on the network segment for a VTN. The mechanisms to identify the network topology or forwarding path with a SID as defined in [RFC8402] are reused, and the control plane can be based on [RFC4915], [RFC5120] and [I-D.ietf-lsr-flex-algo].

2.1. SR-MPLS

For one IGP link, multiple Adj-SIDs are allocated, each of which is associated with a VTN the link participates, and represents a subset of the link resources allocated to the VTN. Similarly, for one IGP node, multiple prefix-SIDs are allocated, each of which is associated with a VTN the node participate, and represent a subset of the node level processing resources allocated to the VTN.

In the case of multi-domain VTNs, on an inter-domain link, multiple BGP peering SIDs [I-D.ietf-idr-bgpls-segment-routing-epe] are allocated, each of which is associated with a VTN which spans

multiple domains, and represents a subset of resources allocated on the inter-domain link.

This way, a group of resource-aware SIDs associated with the same VTN can be used to represent the VTN topology and the allocated resources in data plane. An SR SID-list built with such group of SIDs can be used to steer service traffic to follow a path within the VTN topology, and each SID in the SID-list can also be used to steer the service traffic to be processed with the set of network resources allocated to the VTN.

Note that the introduction of SR-MPLS based VTN increases the number of SIDs needed, and the amount network states is also increased. While thanks to the SR paradigm, the resource-aware SIDs are associated with different VTNs rather than different paths, thus per-path state is still avoided inside the SR network.

2.2. SRv6

In order to support multiple VTNs with SRv6, network nodes (including both the edge nodes and transit nodes) belonging to the same VTN need to have a consistent view of the VTN, and perform consistent computation and forwarding behavior to comply to the VTN topology and resource constraints. The mechanisms to ensure the consistency of such view can be based on [RFC4915], [RFC5120] and [I-D.ietf-lsr-flex-algo]. In data plane, some mechanism is needed for nodes to identify the VTN a packets belongs to.

Based on the mechanisms defined in [I-D.ietf-spring-resource-aware-segments], for a network node, multiple SRv6 LOCs are allocated, each of which is associated with a VTN it participates, and represents a subset of the network resources allocated to the VTN. The SRv6 SIDs of a particular VTN are allocated from the SID space using the VTN-specific LOC as the prefix. These SRv6 SIDs can be used to represent VTN-specific local functions.

A group of SRv6 SIDs associated with the same VTN can be used to represent the VTN topology and the allocated resources in data plane. An SRv6 SID-list built with such SIDs can be used to steer the service traffic to follow a path within the VTN topology, and each SID in the SID-list can also be used to steer the service traffic to be processed with the set of network resources allocated to the VTN.

Note that the introduction of SRv6 based VTN increases the number of SRv6 Locators and SIDs needed, and the amount network states is also increased. While thanks to the SR paradigm, the resource-aware SIDs

are associated with different VTNs rather than different paths, thus per-path state is still avoided in the SR network.

3. Procedures

This section describes the procedures of creating SR based VTNs and the corresponding forwarding tables and entries. Although it is illustrated using SR-MPLS, the proposed mechanism is applicable to both SR-MPLS and SRv6.

According to the received service requirement, a centralized network controller calculates a subset of the network topology to support the service. Within this topology, the set of network resources required on each network element is also determined. The subset of network topology and network resources together constitute a VTN. Depending on the service requirement, the network topology and resource can be dedicated for a individual service or customer, or can be shared by a group of services or customers.

Based on the mechanisms defined in [I-D.ietf-spring-resource-aware-segments], the network topology and resources of a VTN can be represented by a group of resource-aware SIDs. With SR-MPLS, a group of prefix-SIDs and adj-SIDs will be used by network nodes and the network controller to construct an SR based VTN, which is considered as the virtual underlay network for the service. Control plane protocols such as IGP and BGP-LS needs to be extended to distribute the SIDs and the associated resource information of each VTN. The detailed control plane extensions are out of the scope of this document.

Suppose a virtual network is requested by some customer or service. The basic requirement is that customer or service is allocated with some dedicated network resource and does not experience unexpected interference from other services in the same network. Other possible requirements may include the required topology, bandwidth, latency, reliability, etc.

3.1. VTN Topology and Resource Computation

A centralized network controller can be responsible for the planning of a VTN to meet the received service request. The controller needs to collect the information of network connectivity, network resources, network performance and other relevant network states of the underlay network. This can be done using either IGP [RFC5305] [RFC3630] [RFC7471] [RFC8570] or BGP-LS [RFC7752] [RFC8571].

Based on the information collected from the underlay network, the controller obtains the underlay network topology and the information

about the allocated and available network resources. When a service request is received, the controller computes the subset of the network topology, along with the set of the resources needed on each network segment (e.g. links and nodes) in the topology to meet the service requirements, whilst maintaining the needs of the existing services that are using the same network. The subset of network topology and network resources constitute a VTN, which will be used as the virtual underlay network of the requested service.

3.2. VTN Network Resource and SID Allocation

According to the result of VTN planning, the network controller instructs the network nodes with the information of the VTN identifier and the required network resources to be allocated to the VTN, so that the involved network nodes could join the VTN and allocate the network resources accordingly. This can be done with either PCEP [RFC5440] or Netconf/YANG [RFC6241] [RFC7950] with necessary extensions. On each network node involved in the VTN, a set of network resources are allocated on a per virtual network basis, and a group of resource-aware SIDs are allocated to represent the set of resources allocated on the network node and the attached links. Such group of resource-aware SIDs, e.g. prefix-SIDs and adj-SIDs are used as data plane identifiers of the node and links in the VTN.

In underlying forwarding plane, there can be multiple ways of partitioning and allocating a set of network resource to a VTN. For example, [FLEXE] may be used to partition the link resource into different sub-channels to achieve resource isolation between each other. The candidate data plane technologies to support resource partitioning can be found in [I-D.ietf-teas-enhanced-vpn]. The SR SIDs are considered as a unified abstraction in network layer, which can work with various network resource partition and allocation mechanisms in the underlying forwarding plane.

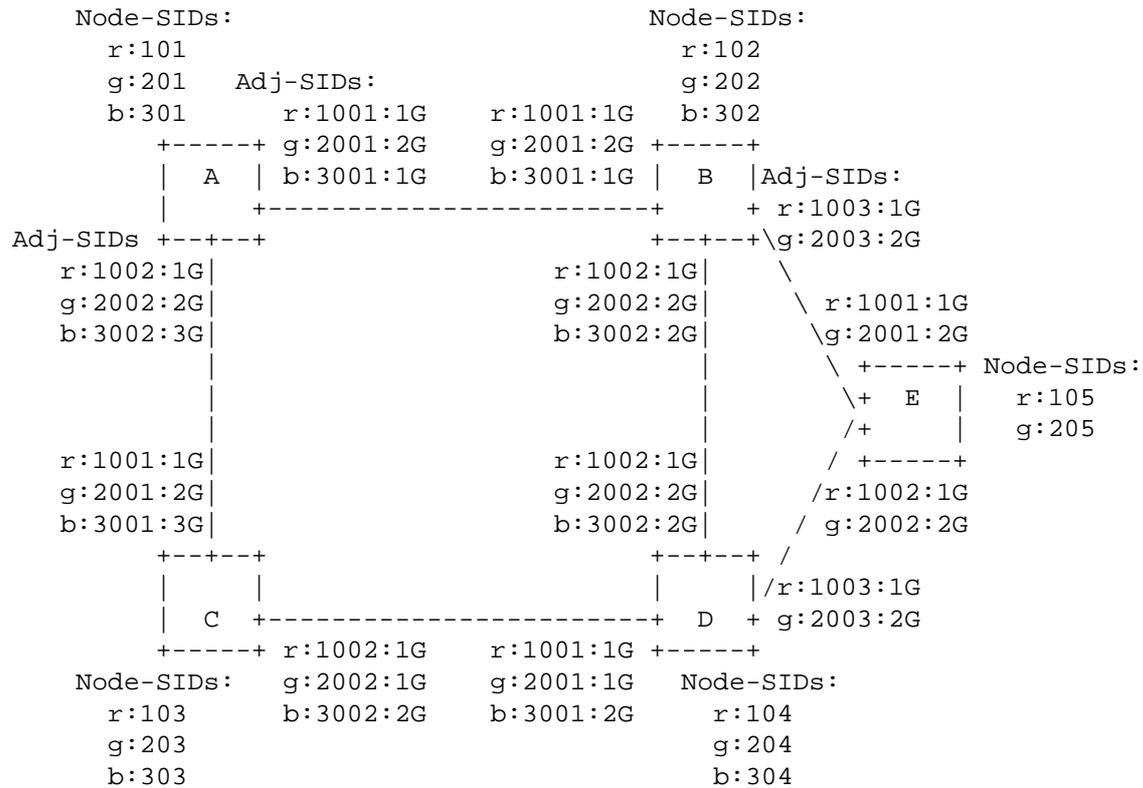


Figure 1. SID and resource allocation for multiple VTNs

Figure 1 shows an example of providing multiple VTNs in an SR based network. Note that the format of the SIDs in this figure is for illustration, both SR-MPLS and SRv6 can be used as the data plane. In this example, three VTNs: red (r) , green (g) and blue (b) are created to carry different services. Both the red and green VTNs consist of nodes A, B, C, D, and E with all their interconnecting links, whilst the blue VTN only consists of nodes A, B, C and D with all their interconnecting links. Note that different VTNs may have a set of shared nodes and links. On each link, a resource-aware adj-SID is allocated for each VTN it participates in.

In Figure 1, the notation x:nnnn:y means that in VTN x, the adj-SID nnnn will steer the packet over a link which has bandwidth y reserved for that VTN. For example, r:1002:1G in link C->D says that the VTN red has a reserved bandwidth of 1Gb/s on link C->D, and will be used by packets arriving at node C with an adj-SID 1002 at the top of the label stack. Similarly, on each node, a resource-aware prefix-SID is allocated for each VTN it participates in. The adj-SIDs can be associated with different set of link resources (e.g. bandwidth) allocated to different VTNs, so that the adj-SIDs can be used to

steer service traffic into different set of link resources in packet forwarding. The prefix-SIDs can be associated with the nodal resources allocated to different VTNs. In addition, the prefix-SIDs can be used to build loose SR path within each VTN, in this case it can be used by the transit nodes to steer service traffic into the set of local network resources allocated to the VTN in forwarding plane.

3.3. Construction of SR based VTNs

In order to make both the network controller and network nodes aware of the information of the VTNs in the network, each network node needs to advertise the identifiers of the VTNs it participates in, together with the group of SIDs and the associated resource attributes both to other network nodes and the controller. This can be achieved by IGP extensions in [I-D.dong-lsr-sr-enhanced-vpn] and BGP-LS extensions in[I-D.dong-idr-bgpls-sr-enhanced-vpn].

Based on the collected information of the topology, the allocated network resource information and the associated SIDs of VTN, the controller and network nodes are able to construct the SR based VTNs and generate the forwarding tables and entries of each VTN based on the prefix-SIDs and adj-SIDs of each VTN. Unlike classic segment routing in which network resources on a network segment are shared by all the SR traffic, different SR VTNs can be associated with different set of resources allocated in the underlay forwarding plane, so that they can be used to meet the enhanced service requirement and provide the required resource isolation from other services in the same network.

Figure 2 shows the SR based VTNs created in the network in Figure 1.

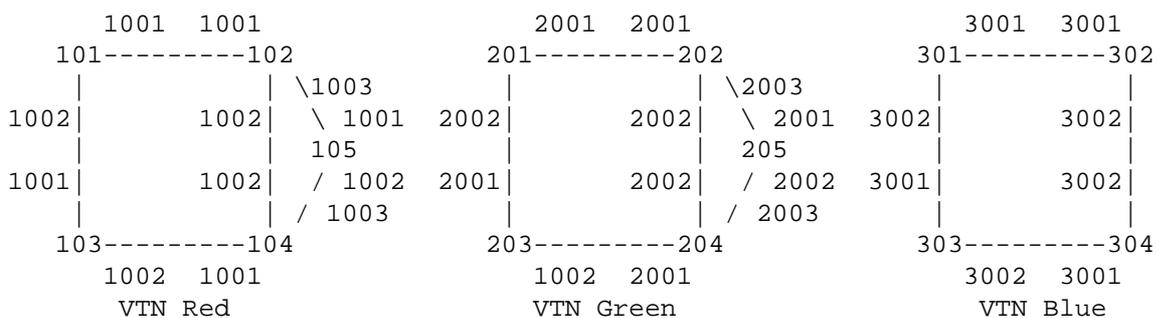


Figure 2. SR based VTNs with different groups of SIDs

For each SR based VTN, SR paths are computed within the VTN, taking the VTN topology and resources as constraints. The SR path can be an explicit path instantiated using SR policy

[[I-D.ietf-spring-segment-routing-policy](#)], in which the SID-list is built only with the SIDs allocated to the VTN. The SR path can also be an IGP computed path associated with a prefix-SID of the VTN, the IGP computation is based on the VTN constraints. Different SR paths in the same VTN may use shared network resources according to the resource-aware SIDs allocated to the VTN, while SR paths in different VTNs can be steered to use different set of network resources over the shared network links or nodes. These VTN-specific SR paths need to be installed in the corresponding forwarding tables.

For example, to create an explicit path A-B-D-E in VTN red in Figure 2, the SR SID-list encapsulated in the service packet would be (1001, 1002, 1003). For the same explicit path A-B-D-E in VTN green, the SR segment list would be (2001, 2002, 2003). In the case where we wish to construct a loose path A-D-E in VTN green, the service packet SHOULD be encapsulated with the SR SID-list (201, 204, 205). At node A, the packet can be sent towards D via either node B or C using the link and node resources allocated for VTN green. At node D the packet is forwarded to E using the link and node resource allocated for VTN green. Similarly, a packet to sent via loose path A-D-E in VTN red would be encapsulated with segment list (101, 104, 105). In the case where an IGP computed path can meet the service requirement, the packet can be simply encapsulated with the prefix-SID of egress node E in the corresponding VTN.

3.4. Mapping Service to SR based VTN

Network services can be provisioned using customized SR based VTN as the underlay network. For example, different services may be provisioned in different SR based VTNs, each of which would use the network resources allocated to the VTN, so that they will not interfere with each other. In another case, a group of services which have similar characteristics and requirement may be provisioned in the same VTN, in this case the network resources allocated to the VTN are only shared among this group of services, but will not be shared with other services in the network. The steering of service traffic to SR based VTNs can be based on either local policy or mechanisms as defined in [[I-D.ietf-spring-segment-routing-policy](#)].

3.5. VTN Visibility to Customer

The customer may request different granularity of visibility to the network which deliver the service. Depending on the requirement, the network can be exposed to the customer either as a virtual network, or a set of computed paths with transit nodes, or simply the abstract connectivity between endpoints without any path information. The visibility can be delivered through different possible mechanisms, such as IGPs (e.g. IS-IS, OSPF) or BGP-LS. In addition, network

operator may want to restrict the visibility of the information it delivers to the customer by either hiding the transit nodes between sites (and only delivering the endpoints connectivity), or by hiding portions of the transit nodes (summarizing the path into fewer nodes). Mechanisms such as BGP-LS allow the flexibility of the advertisement of aggregated virtual network information.

4. Benefits of the Proposed Mechanism

The proposed mechanism provides several key characteristics:

- o **Flexibility:** Multiple customized VTNs can be created in a shared network to meet different customers' connectivity and service requirement. Each customer is only aware of the topology and attributes of his own VTN, and provision services on the VTN instead of the physical network. This provides an efficient mechanism to support network slicing.
- o **Resource Isolation:** The computation and instantiation of SR paths in one VTN can be independent from other VTNs or other services in the network. In addition, a VTN can be associated with a set of network resources, which can avoid resource competition and performance interference from other services in the network. The proposed mechanism also allows resource sharing between different service flows of the same customer, or between a group of services which are provisioned in the same VTN. This gives the operator and the customers the flexibility in network planning and service provisioning. The performance of critical services can be further ensured using the mechanisms defined in [DetNet].
- o **Scalability:** The introduction of resource guaranteed paths or virtual networks would increase the amount of state to the network. The proposed mechanism seeks to achieve a balance between the state limitations of traditional end-to-end TE mechanism and the lack of resource awareness in classic segment routing. Following the segment routing paradigm, network resources are allocated on network segments per VTN and represented as SIDs, thus there is no per-flow state introduced in the network. Operator can choose the granularity of resource allocation to network segments. In network segments where resource is scarce such that the service requirement may not always be met, the proposed approach can be used to allocate specific resources to a VTN which contains such network segment. By contrast, in other segment of the network where resource is considered plentiful, the resource may be shared between a number of VTNs. The decision to do this is in the hands of the operator. Because of the segmented nature of the SR based VTN, resource

aggregation is easier and more flexible than RSVP-TE based approach.

5. Service Assurance

In order to provide service assurance for services provisioned in the SR based VTNs, it is necessary to instrument the network at multiple levels. The network operator needs to ascertain that the underlay network is operating correctly. A tenant needs to ascertain that their services are operating correctly. In principle these can use existing techniques. These are well known problems and solutions either exist or are in development to address them.

New work may be needed to instrument the VTNs that are created for particular services. Such instrumentation needs to operate without causing disruption to other services using the network. Given the sensitivity of some applications, care needs to be taken to ensure that the instrumentation itself does not cause disruption either to the service being instrumented or to other services. In case of failure or performance degradation of a service path in a particular VTN, it is necessary that either local protection or end-to-end protection mechanism is used to switch to another path in the same VTN which could meet the service performance requirement and does not impact other services in the network.

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

The security considerations of segment routing and resource-aware SIDs are applicable to this document.

The SR VTNs may be used carry services with specific SLA parameters. An attack can be directly targeted at the customer application by disrupting the SLA, and can be targeted at the network operator by causing them to violate their SLA, triggering commercial consequences. By rigorously policing ingress traffic and carefully provisioning the resources provided to the VTN, this type of attack can be prevented. However care needs to be taken when shared resources are provided between VTNs at some point in the network, and when the network needs to be reconfigured as part of ongoing maintenance or in response to a failure.

The details of the underlying network should not be exposed to third parties, some abstraction would be needed, this is also to prevent attacks aimed at exploiting a shared resource between VTNs.

8. Contributors

Zhenbin Li
Email: lizhenbin@huawei.com

Zhibo Hu
Email: huzhibo@huawei.com

9. Acknowledgements

The authors would like to thank Mach Chen, Stefano Previdi, Charlie Perkins, Bruno Decraene, Loa Andersson, Alexander Vainshtein and Joel Halpern for the valuable discussion and suggestions to this document.

10. References

10.1. Normative References

[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>>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", [RFC 8660](#), DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

10.2. Informative References

[DetNet] "DetNet WG", 2016, <<https://datatracker.ietf.org/wg/detnet>>.

[FLEXE] "Flex Ethernet Implementation Agreement", March 2016, <<http://www.oiforum.com/wp-content/uploads/OIF-FLEXE-01.0.pdf>>.

[I-D.dong-idr-bgpls-sr-enhanced-vpn]
Dong, J., Hu, Z., Li, Z., Tang, X., and R. Pang, "BGP-LS Extensions for Segment Routing based Enhanced VPN", [draft-dong-idr-bgpls-sr-enhanced-vpn-02](#) (work in progress), June 2020.

- [I-D.dong-lsr-sr-enhanced-vpn]
Dong, J., Hu, Z., Li, Z., Tang, X., Pang, R., JooHeon, L., and S. Bryant, "IGP Extensions for Segment Routing based Enhanced VPN", [draft-dong-lsr-sr-enhanced-vpn-04](#) (work in progress), June 2020.
- [I-D.ietf-idr-bgpls-segment-routing-epe]
Previdi, S., Talaulikar, K., 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-19](#) (work in progress), May 2019.
- [I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-10](#) (work in progress), August 2020.
- [I-D.ietf-spring-resource-aware-segments]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", [draft-ietf-spring-resource-aware-segments-00](#) (work in progress), July 2020.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-08](#) (work in progress), July 2020.
- [I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", [draft-ietf-spring-srv6-network-programming-18](#) (work in progress), August 2020.
- [I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Service", [draft-ietf-teas-enhanced-vpn-06](#) (work in progress), July 2020.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.

- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<https://www.rfc-editor.org/info/rfc3630>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), DOI 10.17487/RFC4915, June 2007, <<https://www.rfc-editor.org/info/rfc4915>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [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>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", [RFC 7471](#), DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.
- [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>>.
- [RFC7950] Bjorklund, M., Ed., "The YANG 1.1 Data Modeling Language", [RFC 7950](#), DOI 10.17487/RFC7950, August 2016, <<https://www.rfc-editor.org/info/rfc7950>>.

- [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>>.

Authors' Addresses

Jie Dong
Huawei Technologies

Email: jie.dong@huawei.com

Stewart Bryant
Futurewei Technologies

Email: stewart.bryant@gmail.com

Takuya Miyasaka
KDDI Corporation

Email: ta-miyasaka@kddi.com

Yongqing Zhu
China Telecom

Email: zhuyq8@chinatelecom.cn

Fengwei Qin
China Mobile

Email: qinfengwei@chinamobile.com

Zhenqiang Li
China Mobile

Email: li_zhenqiang@hotmail.com

Francois Clad
Cisco Systems

Email: fclad@cisco.com