SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: 16 November 2026

J. Dong
Huawei Technologies
T. Miyasaka
KDDI Corporation
Y. Zhu
China Telecom
F. Qin
Z. Li
China Mobile

15 May 2026

Introducing Resource Awareness to SR Segments
draft-ietf-spring-resource-aware-segments-18

Abstract

This document describes a mechanism to allocate network resources to one or a set of Segment Routing Identifiers (SIDs). Such SIDs are referred to as resource-aware SIDs. The resource-aware SIDs retain their original forwarding semantics, with the additional semantics to identify the set of network resources available for the packet processing and forwarding action. The proposed mechanism is applicable to both segment routing with MPLS data plane (SR-MPLS) and segment routing with IPv6 data plane (SRv6).

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 16 November 2026.

Table of Contents

1. Introduction

The Segment Routing (SR) Architecture [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. The base SR specifications do not have the capability of identifying or reserving a set of network resources. Although a centralized controller can have a global view of network state and can provision different services using different SR paths, in data packet forwarding it still relies on the DiffServ QoS mechanism [RFC2474] [RFC2475] to provide coarse-grained traffic differentiation in the network. While such a mechanism may be sufficient for some types of services, others may require a set of dedicated network resources to achieve resource isolation in the same network. Also, the number of such services could be larger than the number of traffic classes available with DiffServ QoS.

Without needing to define new SID types, this document extends the SR paradigm by associating SIDs with network resource attributes, so that network resources can be allocated to one or a set of SIDs. Such SIDs are referred to as resource-aware SIDs. These resource-aware SIDs retain their original functionality, with the additional semantics of identifying the set of network resources available for the packet processing action. Typical types of network resources include link bandwidth, buffers, and queues that are associated with class of service, scheduling weights or time cycles, and it is also possible to associate SR SIDs with other types of resources (e.g., the processing and storage resources). For a particular SR segment, multiple resource-aware SIDs can be allocated, each of which represents a subset of network resources allocated in the network to meet the requirements of one or a group of customers or services. Each subset of the network resources may be associated with one or multiple resource-aware SIDs. The allocation of network resources to segments can be done either via local configuration or via a centralized controller. Other approaches are possible such as the use of a control plane signaling protocol, but they are out of the scope of this document.

An SR Policy that requires dedicated network resources can be composed of a list of resource-aware SIDs. This can be useful for service which requires dedicated network resources along the SR path. In addition, a subset of the network topology and resources (considered as a "virtual network") can be represented by a group of resource-aware SIDs that meet the connectivity and resource goals. The resources associated with each segment of the virtual network can be the same or different. The proposed mechanism is applicable to SR with both MPLS data plane (SR-MPLS) and IPv6 data plane (SRv6). The reader is expected to be familiar with the terminology in [RFC8402], [RFC8660] and [RFC8986].

1.1. 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. Terminology

The following terms are introduced by this document.

* Resource-aware segment: An SR segment which not only represents a specific instruction, but also identifies the set of network resources used for executing the action.

* Resouce group: A group of network resources allocated on a set of network nodes and links, which can be used for forwarding and processing packets with one or multiple resource-aware SIDs.

* Global resource-aware SID: A resource-aware segment identifier which is associated with a resource group.

* Local resource-aware SID: A resource-aware segment identifier which is associated with a specific set of resources on a network node or link in the resource group.

3. Segments with Resource Awareness

In the Segment Routing architecture [RFC8402], several types of segments are defined to represent either topological or service instructions. A topological segment can be a node segment or an adjacency segment. A service segment may be associated with specific service functions for service-chaining purposes. This document introduces additional resource semantics to the existing types of SIDs. A resource-aware SID retains its original functionality, with the additional semantics of identifying a set of network resources allocated in the underlay network for the packet processing action. A resource-aware SID is considered local resource-aware if the associated network resource is allocated on a specific node in the network. A resource-aware SID is considered global resource-aware if the associated network resources are allocated on multiple nodes in the network. A local resource-aware SID may be allocated with a dedicated set of network resources, while for global resource-aware SIDs, a common set of network resources may be allocated to a group of resource-aware SIDs.

This section describes the mechanisms of using resource-aware SR SIDs to indicate the network resource information associated with the SR paths or virtual networks based on the two SR data plane instantiations: SR-MPLS and SRv6. The mechanisms to identify the forwarding path or network topology with SIDs as defined in [RFC8402] do not change. Aligning with the SR architecture, the control plane for resource-aware segments can be centralized, distributed, or hybrid. When resource-aware segments are associated with a virtual network, the control plane for distributing the resource-aware SIDs and the associated topology or Flexible-Algorithm can be based on [RFC4915], [RFC5120] and [RFC9350].

3.1. SR-MPLS

The MPLS instantiation of Segment Routing is specified in [RFC8660]. [RFC8402] specifies several types of SIDs, including IGP Adjacency Segment (Adj-SID), IGP-Prefix Segment (Prefix-SID), and IGP-Node Segment (Node-SID). It also introduces BGP Peer Adjacency Segment (PeerAdj SID). These types of SIDs can be reused to represent both the topological instructions and the set of network resources allocated for packet processing following the instructions.

A resource-aware Adj-SID is a local resource-aware segment, it represents a subset of the network resources (e.g., bandwidth, buffer and queuing resources) on a given link, thus each resource-aware Adj-SID is associated with a subset of the link's traffic engineering (TE) capabilities and resources (known as TE attributes [RFC2702]).

For one IGP link, multiple resource-aware Adj-SIDs can be assigned, each of which is associated with a subset of the link resources allocated on the link. For one inter-domain link, multiple BGP PeerAdj SIDs may be assigned, each of which is associated with a subset of the link resources allocated on the inter-domain link. The inter-domain link is between network domains managed by the same administrative entity and aligns with the trust model described in [RFC8402]. The resource-aware Adj-SIDs may be associated with a specific network topology and/or algorithm, so that it is used only for resource-aware SR paths computed within the topology and/or algorithm.

Note this per-segment resource allocation complies with the SR paradigm, which avoids introducing per-path state into the network. Several approaches can be used to partition and reserve the link resources, such as [FLEXE], logical sub-interfaces with reserved bandwidth, dedicated queues, etc. The detailed mechanism of link resource partitioning is out of scope of this document.

A resource-aware prefix-SID is a global resource-aware segment, it is associated with a network topology and/or an algorithm which the attached node participates in. In addition, a resource-aware prefix-SID is allocated with a set of network resources (e.g., bandwidth, buffer and queuing resources) on the nodes and links participating in the associated topology and/or algorithm. Such set of network resources can be used for forwarding packets which are encapsulated with this resource-aware prefix-SID, along the paths computed in the associated topology and/or algorithm.

Although it is possible that each resource-aware prefix-SID is allocated with a set of dedicated resources on every node and link in the associated topology and/or algorithm, the overhead of per-prefix

resource reservation is usually considered unacceptable in both control plane signaling and data plane states, and it is likely some of the allocated resources will be wasted. It is RECOMMENDED that a common set of network resources be allocated by the network nodes and links participating in the topology and/or algorithm, and this common set of network resources is associated with a group of resource-aware Prefix-SIDs. Such a common set of network resources constitutes a resource group. For a given <topology, algorithm> tuple, there can be one or multiple resource groups. This way, a group of resource-aware prefix-SIDs which are associated with the same <topology, algorithm> tuple can share the set of network resources in a resource group. The association between the SR SIDs and a resource group can be provisioned using the management plane or a control plane.

The recommendation above helps to reduce the dynamics in per-prefix resource allocation and adjustment, so that the network resource can be allocated based on planning and does not have to rely on dynamic signaling. When the set of nodes and links that participate in a <topology, algorithm> tuple changes, the set of network resources allocated on specific nodes and links may need to be adjusted. When the set of network resources are locally configured on the network links, this means that the resources allocated to resource-aware Adj-SIDs on those links may have to be adjusted, and new TE attributes for the associated adj-SIDs re-advertised.

For one IGP prefix, multiple resource-aware prefix-SIDs can be allocated. Each resource-aware prefix-SID may be associated with a unique <topology, algorithm> tuple, in this case different <topology, algorithm> tuples can be used to distinguish the resource-aware prefix-SIDs of the same prefix. In another case, for one IGP prefix, multiple resource-aware prefix-SIDs may be associated with the same <topology, algorithm> tuple but different resource groups. Then an additional control plane distinguisher needs to be introduced to distinguish different resource-aware prefix-SIDs associated with the same <topology, algorithm> but different resource groups. The first approach is simpler and does not require extensions to control plane protocols, while there can be scalability concerns when the number of resource groups is large, as it would require a large number of topologies or Flex-Algorithms. The second approach is more scalable, while it requires additional extensions to the control plane protocols. The exact control plane extensions are out of the scope of this document, but see Section 7 for more discussion of the scalability concerns.

A group of resource-aware Adj-SID and resource-aware Prefix-SIDs can be used to construct the SID lists of an SR Policy, which can be used to steer the traffic to be forwarded along the explicit paths (either strict or loose) and processed using the set of network resources identified by the resource-aware SIDs.

In SR-MPLS packet forwarding, each resource-aware Adj-SID identifies both the next-hop of the node and the set of resources used for packet processing on the outgoing interface. Each resource-aware Prefix-SID identifies the path to the node which the prefix is attached to, and the resource group which consists of a set of network resources to be used for packet forwarding on the transit nodes along the path. The transit nodes use the resource-aware Prefix-SIDs to determine the next-hop of the packet and the set of local resources in the identified resource group, then forward the packet to the next-hop using the set of local resources.

When the set of network resources allocated on the egress node also needs to be determined, it is RECOMMENDED that Penultimate Hop Popping (PHP) [RFC3031] be disabled, otherwise the inner service label needs to be used to infer the set of resources to be used for packet processing on the egress node of the SR path, which would over-complicate the assignment of the service label and potentially require multiple service labels to be assigned for the same service to identify the different resource groups.

This mechanism requires the allocation of additional prefix-SIDs or adj-SIDs to identify different sets of network resources. As the number of resource groups increases, the number of SIDs would increase accordingly, while it should be noted that there is still no per-path state introduced into the network.

3.2. SRv6

[RFC8986] defines the SRv6 SID format (LOC:FUNCT:ARG) and the base set of SRv6 behaviors bound to the SRv6 SIDs. When the LOC (Locator) part of the SRv6 SIDs is routable, it leads to the node which instantiates the SID.

The approach of introducing resource-awareness to SRv6 is by firstly making the SRv6 Locators resource-aware. For one SRv6 node, multiple resource-aware SRv6 Locators can be assigned. A resource-aware Locator is associated with a network topology and/or algorithm in which the originating node participates, as well as a set of network resources (e.g., bandwidth, buffer, and queueing resources) on each node and the attached links participating in the same topology and/or algorithm. Then resource-aware SRv6 SIDs are allocated using the resource-aware SRv6 Locator as the prefix, and the resource-aware

SRv6 SIDs are associated with a subset of the local resources which belong to the resource group associated with the resource-aware SRv6 Locator. The set of network resources allocated to the resource-aware SRv6 Locators and SRv6 SIDs are used for forwarding packets in which the resource-aware SRv6 SIDs are encoded as the destination IPv6 addresses.

Similar to the approach used with resource-aware prefix-SIDs in SR-MPLS, it is RECOMMENDED that a common set of network resources are allocated by the network nodes and links participating in a topology and/or algorithm, and this resource group is associated with a group of resource-aware Locators of the same topology and/or algorithm.

For one IGP link, multiple resource-aware SRv6 End.X SIDs can be allocated to identify different sets of link resources allocated on the link. SRv6 SIDs for other types of behaviors MAY also be assigned as resource-aware SIDs, which identifies the set of network resources allocated by the node for executing the behavior. All resource-aware SRv6 SIDs MUST use a resource-aware locator as its prefix.

A group of resource-aware SRv6 SIDs can be used to construct the SID lists of an SR Policy, which can be used to steer the traffic to be forwarded along the explicit paths (either strict or loose), and be processed using the set of network resources identified by the resource-aware SRv6 Locators and SIDs.

In SRv6 packet forwarding, the transit nodes use the resource-aware Locator of the SRv6 SID carried in the destination IPv6 address field to determine the next-hop of the packet, and the resource group which consists of a set of network resources to be used for packet forwarding along the path. On the segment endpoint nodes, the resource-aware End.X SID identifies both the next-hop and the set of resources used for packet forwarding on the outgoing interface of the node which instantiates the SID.

This mechanism requires the allocation of additional SRv6 Locators and SIDs to identify different set of network resources. As the number of resource groups increases, the number of SRv6 Locators and SIDs would increase accordingly, while it should be noted that there is still no per-path state introduced into the network.

4. Control Plane Considerations

The mechanism described in this document assumes the use of a centralized controller to collect the information about the network (configuration, state, routing databases, etc.) as well as the service information (traffic matrix, performance statistics, etc.) for the planning of network resources based on the service requirements. The centralized controller can also be used to instruct the network nodes to allocate the network resources and associate the resources to resource-aware SIDs. The resource-aware SIDs can be either explicitly provisioned by the controller, or can be dynamically allocated by network nodes. The distributed control plane is complementary to the centralized controller. When the resource-aware SIDs are locally configured or dynamically allocated, a distributed control plane can be used for the collection and distribution of the resource-aware SIDs among network nodes, together with the set of associated local network resource information. Then some of the network nodes can distribute the collected information to the centralized controller. The mechanisms as defined in [RFC8665][RFC8667] [RFC9085] [RFC9352] [RFC9513] and [RFC9514] can be reused with possible extensions to improve the efficiency and scalability. It is anticipated that augmentations to the SR YANG models would provide a way to configure and learn about the relationship between resource-aware SIDs and the resources on specific nodes.The details are out of the scope of this document. If an attempt to associate a resource-aware SID with resources on a router fails (for example, due to an error in the amount of resource requested) the control plane solution MUST report this so that the situation can be corrected.

The support for a resource group and the SR SIDs or SRv6 locators information to associate packets to it MUST be aligned among the network nodes in that resource group, so as to ensure that packets are processed consistently within a resource group. This task can be accomplished via local configuration or via a centralized controller. Other approaches may be possible. [I-D.ietf-teas-nrp-yang] provides some guidance on the provisioning of resource-aware segments for network resource partitions (NRPs). The centralized controller or management system is responsible for consistent provisioning of resource groups, and should be able to roll back in case of partial provisioning failure. A resource group SHOULD NOT be used until it is fully provisioned. An update to a resource group is finished until all changes to the involved network nodes are successfully made.

To indicate the support for a given resource group, a node needs to advertise the identifier of the resource group, the associated topology and algorithm, the resource-aware SIDs and the TE attributes

representing the resources allocated to the resource-aware SIDs in the resource-group. The TE attributes of a resource group may be used as constraints in path computation.

The controller is also responsible for the centralized computation and optimization of the SR paths taking the topology, algorithm and network resource constraints into consideration. The interaction between the controller and the network nodes can be based on Netconf/YANG [RFC6241] [RFC7950] [I-D.ietf-spring-sr-policy-yang], BGP SR Policy [RFC9830] or PCEP [RFC8664] [RFC9603]. In some scenarios, extensions to some of these protocols may be needed to improve the efficiency and scalability of the control plane, but these are out of the scope of this document. Distributed computation of resource-aware SR paths is also possible, the topology, algorithm and/or resource constraints need to be taken into consideration by network nodes. The distributed control plane may be based on [RFC4915], [RFC5120], [RFC9350] with necessary extensions.

When a network node is instructed to associate a SID with specific resources, its actions will depend on the operational mechanisms of the network. In some cases the association between SIDs and resources is configured on the individual network nodes, and the control plane (e.g. IGP) is used to distribute the SID information and the allocated resource information to the controller and the ingress nodes for TE constraint-based path computation. In network cases with SR and other TE mechanisms (such as RSVP-TE) co-existing in the network, the IGP advertisements of available resources may need to be updated to indicate that there has been a change to the available resources resulting from the instantiation of a new resource-aware SID, it is suggested such updates would be rate-limited to avoid overloading the IGP system using suppression mechanisms as described in [RFC8570] [RFC7471]. In still other cases the association between SIDs and network resources is provisioned by the central controller which is responsible for all TE management, then the distributed control plane does not need to take any additional action.

5. IANA Considerations

This document makes no request of IANA.

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

6. Implementation Status

This section is to be removed before publishing as an RFC.

RFC-Editor: Please clean up the references cited by this section before publication.

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

According to [RFC7942], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

This section is provided in compliance with the SPRING working group policies ([SPRING-WG-POLICIES]).

6.1. Huawei Technologies

Huawei Technologies reported the following implementations of the resource-aware segments (Section 2). The resource-aware segments are used to build SR based Network Resource Partitions (NRPs) and resource guaranteed SR Policies.

* Huawei ATN9XX, CX600 routers.

* Huawei NE40E, NE8000, NE5000E routers.

At the time of this report, all the implementations listed above are in production and follow the specification in the latest version of this document, including all the "MUST" and "SHOULD" clauses for the resource-aware segments.

This report was last updated on August 28, 2025.

7. Operational Considerations

Resource-aware segments can coexist with the existing SR segments. Network operators may introduce resource-aware segments into a portion of their SR networks to support services which require guaranteed network resources (e.g. bandwidth). The use of either base SR segments or resource-aware SR segments for specific service is based on operators' local policy.

Resource-aware segments require to introduce additional SR-MPLS SIDs or SRv6 Locators/SIDs for different subsets of network resources. This would increase the amount of SR SIDs to be managed, and would also increase the amount of state to be maintained by network nodes. Although with the SR paradigm, per-path state can be avoided in the network. The scalability of the deployed solution may also depend on the control plane solution that is available in implementations. If no additional control plane features are available, the only choice is to use different <topology, algorithm> tuples to distinguish the resource-aware prefix-SIDs of the same prefix. This approach may be suitable for small numbers of resource groups (less than ten or so), but with more resource groups, this approach will require more topologies or Flex-Algorithms, each of which requires separate management and can stress operational systems. If a larger number of resource groups are required, then operators should use the alternate method to allocate additional prefix-SIDs or adj-SIDs to identify the resource groups, but must utilize additional control plane mechanisms (see Section 3) to distribute the association of SIDs to resource groups. Operators need to be aware of the additional cost of introducing resource-aware segments, and provide careful planning of the resource groups, so that the resource-aware segments can meet the service requirements without introducing unacceptable complexity to network operation and management.

The consistency in the binding between resource-aware segments and resource groups across all participating nodes in the network is crucial for correct and consistent treatment to packets so as to meet the resource guarantee and SLA requirements. If this is not the case, it may cause problems including service quality degradation or packet drop. Such issues could be detected and diagonosed using performance measurement or packet trace mechanisms with the same resource-aware segments as in the data packets used for forwarding. Control plane mechanisms need to include consistency checks to allow the configured state of resource allocation in network nodes to be verified against the intended state. If inconsistency in resource binding is detected by a network node, by default the impacted resource-aware SIDs MUST NOT be used for traffic forwarding, and an error SHOULD be logged and reported.

Operators need to specify the policy for traffic which exceeds the allocated resources of the resource-aware SID carried. The options include: drop the traffic, lower the priority and treat it in best effort, etc.

8. Security Considerations

The security considerations of segment routing and SRv6 in [RFC8402] [RFC8660] [RFC8754], [RFC8986] and [I-D.ietf-spring-srv6-security] are applicable to this document.

The allocaton of network resources, the association of resource-aware SIDs with the allocated network resources, and the distribution of information of the resource-aware SIDs together with the associated TE attributes MUST be done via control or management protocol channels with proper mechanisms for authentication, authorization, integrity or replay-protection. The specifications of the control or managment plane protocols for resource-aware segments SHOULD specify how these security properties are provided. When the control plane of resource-aware segments is based on Flex-Algo, the security threats described in [RFC9350] need to be considered, as the hijack of a Flex-Algo which associates with an resource group would compromise not just path selection but also resource isolation correctness.

A compromised or misconfigured controller, or a node with local configuration authority, could allocate sufficient network resources and resource-aware SIDs to exhaust link or node resources, thereby starving the base SR forwarding plane. The allocation of network resources and resource-aware SIDs MUST be under some admission control, and implementations SHOULD reject network resource and resource-aware SIDs allocation when it would exceed a configurable threshold, ensuring that base SR forwarding plane availability cannot be compromised by resource exhaustion.

The resource-aware SIDs may be used for provisioning of SR paths or virtual networks to carry traffic with specific SLA requirements (such as latency). By disrupting the SLA of such traffic an attack can be directly targeted at the customer application, or can be targeted at the network operator by causing them to violate their SLA, triggering commercial consequences. Dynamic attacks of this sort are not something that networks have traditionally guarded against, and networking techniques need to be developed to defend against this type of attack. By rigorously policing ingress traffic and carefully provisioning network resources provided to such services, this type of attack can be prevented. However care needs to be taken when providing shared resources, and when the network needs to be reconfigured as part of ongoing maintenance or in response to a failure.

A compromised network node may choose not to actually allocate the claimed resources to the resource-aware SIDs, overstate the available resources, or selectively degrade specific resource groups, this may result in the expected SLA being disrupted due to lack of resource guarantee.

The resource-aware SIDs carried in data packets can reveal not just where the packets go, but also the corresponding resource groups. The details about resource allocation in the underlay network MUST NOT be exposed to third parties, so as to prevent attacks aimed at exploiting shared network resources.

9. Contributors

Stewart Bryant
Email: stewart.bryant@gmail.com

Francois Clad
Email: fclad@cisco.com

Zhenbin Li
Email: lizhenbin@huawei.com

Zhibo Hu
Email: huzhibo@huawei.com

Joel Halpern
Email: jmh@joelhalpern.com

10. Acknowledgements

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

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>.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <https://www.rfc-editor.org/info/rfc3031>.

[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>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

11.2. Informative References

[FLEXE] "Flex Ethernet Implementation Agreement", March 2016, <https://www.oiforum.com/wp-content/uploads/2019/01/OIF-FLEXE-01.0.pdf>.

[I-D.ietf-spring-sr-policy-yang] Saleh, T., Raza, S. K., Zhuang, S., Matsushima, S., and V. P. Beeram, "YANG Data Model for Segment Routing Policy", Work in Progress, Internet-Draft, draft-ietf-spring-sr-policy-yang-06, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-policy-yang-06>.

[I-D.ietf-spring-srv6-security] Buraglio, N., Mizrahi, T., tongtian124, Contreras, L. M., and F. Gont, "Segment Routing IPv6 Security Considerations", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-security-14, 13 April 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-security-14>.

[I-D.ietf-teas-nrp-yang] Wu, B., Dhody, D., Beeram, V. P., Saad, T., and S. Peng, "YANG Data Models for Network Resource Partitions (NRPs)", Work in Progress, Internet-Draft, draft-ietf-teas-nrp-yang-05, 22 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-teas-nrp-yang-05>.

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <https://www.rfc-editor.org/info/rfc2474>.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, DOI 10.17487/RFC2475, December 1998, <https://www.rfc-editor.org/info/rfc2475>.

[RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, DOI 10.17487/RFC2702, September 1999, <https://www.rfc-editor.org/info/rfc2702>.

[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>.

[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>.

[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>.

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://www.rfc-editor.org/info/rfc7942>.

[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>.

[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>.

[RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <https://www.rfc-editor.org/info/rfc8665>.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <https://www.rfc-editor.org/info/rfc8667>.

[RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <https://www.rfc-editor.org/info/rfc9085>.

[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>.

[RFC9087] Filsfils, C., Ed., Previdi, S., Dawra, G., Ed., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", RFC 9087, DOI 10.17487/RFC9087, August 2021, <https://www.rfc-editor.org/info/rfc9087>.

[RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <https://www.rfc-editor.org/info/rfc9350>.

[RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <https://www.rfc-editor.org/info/rfc9352>.

[RFC9513] Li, Z., Hu, Z., Talaulikar, K., Ed., and P. Psenak, "OSPFv3 Extensions for Segment Routing over IPv6 (SRv6)", RFC 9513, DOI 10.17487/RFC9513, December 2023, <https://www.rfc-editor.org/info/rfc9513>.

[RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <https://www.rfc-editor.org/info/rfc9514>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <https://www.rfc-editor.org/info/rfc9552>.

[RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <https://www.rfc-editor.org/info/rfc9603>.

[RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <https://www.rfc-editor.org/info/rfc9830>.

[SPRING-WG-POLICIES] Chairs, S. W. G., "SPRING Working Group Policies", 14 October 2022, <https://wiki.ietf.org/en/group/spring/WG_Policies>.

Authors' Addresses

Jie Dong
Huawei Technologies
Email: jie.dong@huawei.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