SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 18, 2020

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

December 16, 2019

Segment Routing for Resource Partitioned Virtual Networks
draft-dong-spring-sr-for-enhanced-vpn-06

Abstract

This document describes the mechanism to associate Segment Routing Identifiers (SIDs) with network resource attributes. The resource-aware SIDs retain their original functionality, with the additional semantics of identifying the set of network resources available for the packet processing action. These SIDs can be used to build SID lists with reserved network resources, which can be used for many network scenarios. One typical use case is to provide SR virtual networks with required network topology and resource attributes. The proposed mechanism is applicable to both segment routing with MPLS data plane (SR-MPLS) and segment routing with IPv6 data plane (SRv6).

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of This Memo

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 June 18, 2020.

Copyright Notice

Copyright (c) 2019 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

Segment Routing (SR) [RFC8402] specifies a mechanism to steer packets through an ordered list of segments. A segment is often referred to by its Segment Identifier (SID). With SR, explicit source routing can be achieved without introducing per-path state into the network. 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 particular services or customers. Although a centralized controller can have a global view of network state and can provision different services onto different SR paths, in packet forwarding it still relies on traditional DiffServ QoS mechanism [RFC2474] [RFC2475] to provide coarse-grained traffic differentiation in the network. While such kind of mechanism may be sufficient for some types of services, it may not meet the stringent requirement of some enhanced services which require dedicated network resources to achieve isolation from other services in the network. Also the number of such enhanced services can be larger than the number of classes in DiffServ QoS.

This document extends the SR paradigm by associating SIDs with network resource attributes. 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. For a particular network segment, multiple resource-aware SIDs can be allocated, each of which represents different set of network resources allocated to meet different service requirement. This mechanism is applicable to SR with both MPLS data plane (SR-MPLS) and IPv6 data plane (SRv6).

The proposed resource-aware SIDs can be used to build SID lists with reserved network resources, which can be used in network scenarios which require to allocate network resources to the processing of particular service traffic. One typical use case is to provide SR based virtual networks with required network topology and resource attributes. A group of resource-aware SIDs 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 corresponding set of network resources. In this case the proposed mechanism can provide the underlay for enhanced VPN services as described in [I-D.ietf-teas-enhanced-vpn].

2. Segment Routing with Topology and Resource Awareness

When SR is used as the data plane to provide different virtual networks in the same network, it is necessary that the SR paths are computed within a virtual network topology, and are instantiated with the corresponding set of network resources.

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 function for service chaining purpose. This document introduces additional resource semantics to SIDs, so that the SIDs can be used to identify a particular network topology and the set of network resources.

This section describes the mechanisms to identify the virtual network topology and resource information with the two SR data plane instantiations: SR-MPLS and SRv6.

2.1. SR-MPLS

In SR-MPLS [I-D.ietf-spring-segment-routing-mpls], IGP Adjacency Segment (Adj-SID) is an IGP-segment attached to a unidirectional adjacency or a set of unidirectional adjacencies. IGP Prefix segment is an IGP segment attached to an IGP prefix, and IGP node segment is an IGP segment that identifies a specific router (e.g., a loopback). In [I-D.ietf-idr-bgpls-segment-routing-epe], PeerAdj SID is used as instruction to steer over a specific local interface towards a specific peer node in a peering Autonomous System (AS). These types of SIDs can be extended to represent both topological elements and the resources allocated on a network element.

For one IGP link, multiple Adj-SIDs SHOULD be allocated, each of which is associated with a virtual network topology, and MAY represent a subset of link resources. Several approaches can be used

to partition the link resource, such as [FLEXE], layer-2 logical sub-interfaces, dedicated queues, etc. The detailed mechanism of resource partitioning is out of scope of this document. Similarly, for one IGP node, multiple prefix-SIDs SHOULD be allocated, each of which is associated with a virtual network topology, and may represent a subset of the node resource (e.g. the processing resources). For one inter-domain link, multiple BGP PeerAdj SIDs [I-D.ietf-idr-bgpls-segment-routing-epe] SHOULD be allocated, each of which is associated with a specific virtual network topology which spans multiple domains, and MAY represent a subset of link resource on the inter-domain link. Note that this per-segment resource allocation complies to the SR paradigm, which avoids introducing per-path state into the network.

A group of SIDs associated with the same virtual network can be used to construct the SR SID-lists (either strict or loose) to steer the traffic of a particular service within the virtual network. Each SID in the SID-list MAY also represent the set of network resources reserved on a network segment.

In data packet forwarding, the SIDs are used to identify the virtual network the packet belongs to, so that a virtual network specific next-hop can be determined. The adj-SIDs MAY also be used to steer traffic of different services into different set of link resources. The prefix-SIDs MAY be used to steer traffic of different services into different set of node resources. When a prefix-SID is used in the SID-list to build an SR loose path, the transit nodes use the prefix-SID to identify the virtual network, and MAY process the packet using the local resources allocated for the corresponding virtual network. Note in this case, it is RECOMMENDED that Penultimate Hop Popping (PHP) [RFC3031] be disabled, otherwise the inner service label SHOULD be used to infer the set of resources to be used.

This mechanism requires to allocate additional prefix-SIDs and adj-SIDs for each virtual network. As the number of virtual networks increases, the number of SIDs would increase accordingly. It is expected that this mechanism is applicable to networks with a limited number of virtual networks.

2.2. SRv6

As specified in [I-D.ietf-spring-srv6-network-programming], an SRv6 Segment Identifier (SID) is a 128-bit value which consists of a locator (LOC) and a function (FUNCT), optionally it may also contain additional arguments (ARG) immediately after the FUNCT. The LOC of the SID is routable and leads to the node which instantiates that SID, which means the LOC can be parsed by all nodes in the network. The FUNCT part of the SID is an opaque identification of a local function bound to the SID, which means the FUNCT and ARG parts can only be parsed by the node which instantiates that SID.

In order to support multiple virtual networks in a SRv6 network, all the nodes (including the edge nodes and transit nodes) belonging to the same virtual network MUST have a consistent view of the virtual network, and performs consistent computation and forwarding behavior to comply to the network topology and resource constraints. A node which participates in multiple virtual networks MUST be able to distinguish packets which belong to different virtual networks.

Taking the above into consideration, for a network node, multiple SRv6 LOCs SHOULD be allocated, each of which is associated with a virtual network topology, and MAY represent a subset of the network resources associated with the virtual network. The SRv6 SIDs of a particular virtual network SHOULD be allocated from the SID space using the virtual network specific LOC as the prefix. These SRv6 SIDs can be used to represent virtual network specific local functions.

A group of SRv6 SIDs associated with the same virtual network can be used to construct the SR SID-lists (either strict or loose) to steer the traffic of a particular service within the virtual network. Each SID in the SID-list MAY also represent the set of network resources which are reserved on a network segment.

In data packet forwarding, the LOC part of SRv6 SID is used by transit nodes to identify the virtual network the packet belongs to, so that a virtual network specific next-hop can be determined. The LOC MAY also be used to indicate the set of local network resources on the transit nodes to be used for the forwarding of the received packet. The SRv6 segment endpoint nodes use the virtual network specific SRv6 SID to identify the virtual network the packet belongs to, and the particular local function to perform on the received packet. The local SRv6 SID MAY also be used to identify the set of

network resource to be used for executing the local function.

This mechanism requires to allocate additional SRv6 Locators and SIDs for each virtual network. As the number of virtual networks increases, the number of Locators and SIDs would increase accordingly. It is expected that this mechanism is applicable to networks with a limited number of virtual networks.

3. Control Plane Considerations

The mechanism described in this document makes 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 virtual networks. The controller is also responsible for the centralized computation and optimization of the SR paths within different virtual networks. The SR SIDs can be either explicitly provisioned by the controller, or dynamically allocated by network nodes then reported to the controller. The interaction between the controller and the network nodes can be based on PCEP [RFC5440], Netconf/YANG [RFC6241] [RFC7950] and BGP-LS [RFC7752]. In some scenarios, extensions to some of these protocols is needed, which are out of the scope of this document and will be specified in separate documents. In some cases a centralized controller may not be used, but this would complicate the operations and planning therefore not suggested.

A distributed control plane can be used for the collection and distribution of the network topology and resource information of the virtual networks among network nodes. Distributed route computation for services within a virtual network is also needed. The distributed control plane is complementary to the centralized controller.

4. Procedures

This section describes the procedures of creating SR based virtual networks and the corresponding forwarding tables and entries.

According to the received service requirement, a centralized network controller calculates a subset of the physical network topology to support the service. Within this topology, the set of network

resources required on each network element is also determined. This network topology and the set of network resources together constitute a virtual network. Depending on the service requirement, the network topology and resource can be dedicated for a particular service, or can be shared with other services.

Following the segment routing paradigm, the network topology and resource are represented using a group of dedicated SIDs. The group of prefix-SIDs and adj-SIDs allocated for a virtual network will be used by network nodes and the network controller to construct an SR based virtual network, which is considered as the underlay network for the service. IGP and BGP-LS needs to be extended to distribute the SIDs and the associated resource information of each virtual network. The IGP extensions are specified in [I-D.dong-lsr-sr-enhanced-vpn], and the BGP-LS extensions are specified in [I-D.dong-idr-bgpls-sr-enhanced-vpn].

Suppose tenant A requests for a virtual network from the network operator. The requirement is that services of tenant A has dedicated network resource allocated and does not experience unexpected interference from other services in the same network, such as other tenants' VPN services, or non-VPN services in the network. The detailed requirements can be described with characteristics such as the following:

- Service topology: the service sites and the connectivity between them.
- Service bandwidth: the bandwidth requirement between service sites.
- Isolation: the level of isolation from other services in the network.
- Reliability: whether fast local repair or end-to-end protection is needed or not.
- Latency: the maximum latency between specific service sites.

- Visibility: the customer may want to have some form of visibility

of the virtual network delivering the service.

4.1. Virtual Network Topology and Resource Computation

As described in section 4, a centralized network controller is responsible for the planning of a virtual network to meet the received service request. The controller collects 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] [RFC7810] 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 from a tenant, the controller computes the subset of the network topology, along with the set of the resources needed on each network element (e.g. links and nodes) in the topology to meet the tenant's service requirements, whilst maintaining the needs of the existing tenants that are using the same network. The subset of network topology and resource constitute a virtual network, which will be used as the underlay of the requested service.

4.2. Network Resource and SID Allocation

According to the result of virtual network planning, network controller instructs the network devices involved in the virtual network to join the virtual network and allocate the required network resources for the virtual network. This can be done with either PCEP [RFC5440] or Netconf/YANG [RFC6241] [RFC7950] with necessary extensions. The network resources are allocated in a per-segment manner. In addition, a set of dedicated SIDs, e.g. prefix-SIDs and adj-SIDs are allocated to represent the virtual network and the network resources allocated on each network segment for this virtual network.

In the underlay forwarding plane, there can be multiple ways of partitioning and allocating a set of network resource to a virtual network. For example, [FLEXE] may be used to partition the link resource into different sub-channels to achieve hard 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 used as a unified abstraction in network layer for various network resource partition and allocation mechanisms in the underlying forwarding plane.

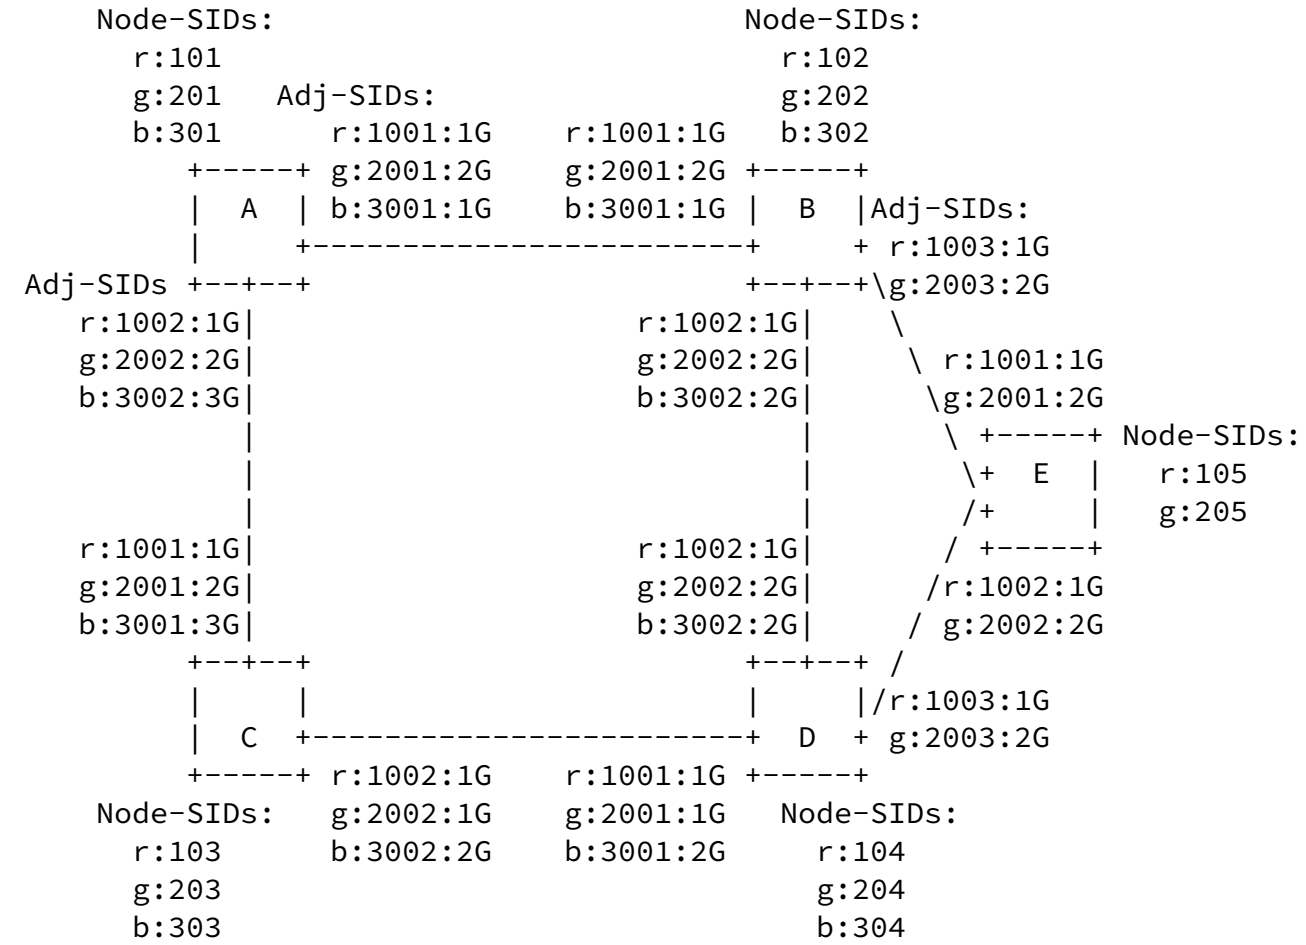

Figure 1. SID and resource allocation for multiple virtual networks

Figure 1 shows an exmple of SR network to support multiple virtual networks. 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 virtual networks: red (r) , green (g) and blue (b) are created to carry different services. Both the red and green virtual networks consist of nodes A, B, C, D, and E with all their interconnecting links, whilst the blue virtual network only consists of nodes A, B, C and D with all their interconnecting links. Note that different virtual networks may have a set of shared nodes and links. On each link, a dedicated adj-SID is allocated for each virtual network it participates in. In Figure 1, the notation x:nnnn:y means that in virtual network x, the adj-SID nnnn will steer the packet over a link which has bandwidth y reserved for that virtual network. For example, r:1002:1G in link C->D says that the virtual network 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 dedicated prefix-SID is allocated for each virtual network it participates in.

The adj-SIDs can be associated with different set of link resources (e.g. bandwidth) allocated to different virtual networks, 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 virtual network. In addition, the prefix-SIDs can be used to build loose SR path within each virtual network, in this case it can be used by the transit nodes to steer different service traffic into different set of local network resources in the forwarding plane.

4.3. Construction of SR based Virtual Networks

In order to make both the network controller and network nodes aware of the information of the virtual networks in the network, each network node SHOULD advertise the identifiers of the virtual networks it participates in, together with the group of SIDs and the associated resource attributes both to other nodes in the network and to 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 virtual network topology, the associated network resource and SIDs information, the controller and network nodes are able to construct the SR virtual network and generate the forwarding tables and entries of each virtual network based on the prefix-SIDs and adj-SIDs allocated for each virtual network. Unlike classic segment routing in which network resources are shared by all the services, different SR virtual networks can be associated with different set of resource allocated in the underlay forwarding plane, so that they can be used to meet the enhanced service requirement and provide the required isolation from other services in the same network.

Figure 2 shows the SR based virtual networks created in the network in Figure 1.

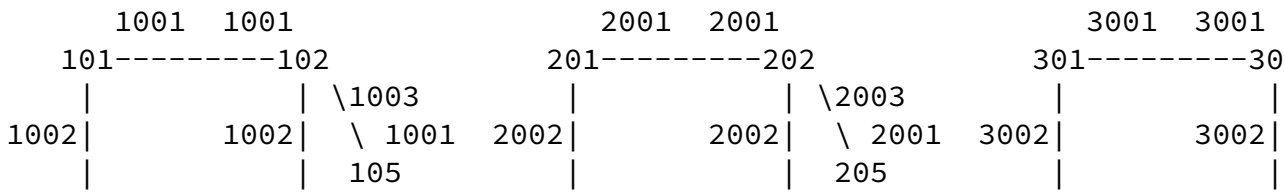

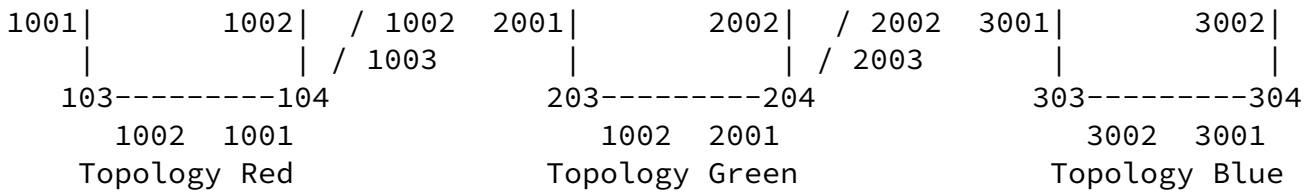

Figure 2. SR based virtual networks with different groups of SIDs

For each SR virtual network, SR paths are computed within the virtual network, taking its network 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 virtual network. The SR path can also be an IGP computed path associated with a particular prefix-SID of the virtual network. Different SR paths in the same virtual network would share the network resources allocated to the virtual network, while SR paths in different virtual networks can be steered to use different set of network resources on the shared network links or nodes.

For example, to create an explicit path A-B-D-E in virtual network 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 virtual network 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 virtual network 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 virtual network green. At node D the packet is forwarded to E using the link and node resource allocated for virtual network green. Similarly, a packet to sent via loose path A-D-E in virtual network 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 node SID of egress node E in the corresponding virtual network.

4.4. Service to SR Virtual Network Mapping

Network services can be provisioned using customized SR virtual networks as the underlay network. For example, different services

may be provisioned in different SR virtual networks, each of which would use the network resources allocated to a particular virtual network, so that they will not interfere with each other. In another case, a group of services which have similar characteristics and requirement can be provisioned in the same SR virtual network, in this case the network resources allocated to the virtual network are shared by these set of services, but will not be shared with other services in the network.

4.5. Virtual Network Visibility to Customer

The tenants of service may request different granularity of visibility to the network which deliver the service. Depending on the requirement, the network can be exposed to the tenant 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 tenant 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.

5. Benefits of the Proposed Mechanism

The proposed mechanism provides several key characteristics:

- Flexibility: Multiple customized virtual networks can be created in a shared network to meet different tenants' connectivity and service requirement. Each tenant is only aware of the topology and attributes of his own virtual network, and provision services on the virtual network instead of the physical network. This provides an efficient mechanism to support network slicing.

- Isolation: Each virtual network can have independent SR path computation and instantiation. In addition, a virtual network 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 services in the same virtual network, or between a group of services which are provisioned in different virtual networks. This gives the operator and the tenants the flexibility in network planning and service provisioning. The performance of critical services can be further ensured using the mechanisms defined in [DetNet].

- Scalability: 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 basic segment routing. Following the segment routing paradigm, network resources are allocated on network segments 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 the segment in a virtual network. By contrast, in other segment of the network where resource is considered plentiful, the resource may be shared between a number of virtual networks. The decision to do this is in the hands of the operator. Because of the segmented nature of the virtual network, resource aggregation is easier and more flexible than RSVP-TE based approach.

6. Service Assurance

In order to provide service assurance for services provisioned in the SR virtual networks, 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 is needed to instrument the virtual networks 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 virtual network, it is necessary that either

local protection or end-to-end protection mechanism is used to switch to another path in the same virtual network which could meet the service performance requirement and does not impact other services in the network.

7. IANA Considerations

This document makes no request of IANA.

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

8. Security Considerations

The security considerations of segment routing are applicable to this document.

The Resource-aware SIDs may be used for provisioning of SR paths or virtual networks to carry traffic with latency SLAs. By disrupting the latency 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 the 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.

The details of the underlay network MUST NOT be exposed to third parties, to prevent attacks aimed at exploiting a shared resource.

9. Contributors

The following people contributed to the content of this document.

10. Acknowledgements

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

11. References

11.1. Normative References

[I-D.ietf-spring-segment-routing-mpls]
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-22 (work in progress), May 2019.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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>.

11.2. 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. and Z. Hu, "BGP-LS Extensions for Segment Routing based Enhanced VPN", draft-dong-idr-bgpls-sr-enhanced-vpn-00 (work in progress), November 2019.

[I-D.dong-lsr-sr-enhanced-vpn]
Dong, J., Hu, Z., and S. Bryant, "IGP Extensions for Segment Routing based Enhanced VPN", draft-dong-lsr-sr-enhanced-vpn-02 (work in progress), November 2019.

[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-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-06 (work in progress), December 2019.

[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-06 (work in progress), December 2019.

[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-03 (work in progress), September 2019.

[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>.

[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>.

[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>.

[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>.

[RFC5439] Yasukawa, S., Farrel, A., and O. Komolafe, "An Analysis of Scaling Issues in MPLS-TE Core Networks", RFC 5439, DOI 10.17487/RFC5439, February 2009, <https://www.rfc-editor.org/info/rfc5439>.

[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>.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <https://www.rfc-editor.org/info/rfc6790>.

[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>.

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

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[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: zhuyq.gd@chinatelecom.cn

Fengwei Qin
China Mobile

Email: qinfengwei@chinamobile.com

Zhenqiang Li
China Mobile

Email: li_zhenqiang@hotmail.com