Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2019

J. Dong
S. Bryant
Huawei Technologies
Z. Li
China Mobile
T. Miyasaka
KDDI Corporation

March 11, 2019

Segment Routing for Enhanced VPN Service
draft-dong-spring-sr-for-enhanced-vpn-03

Abstract

Enhanced VPN (VPN+) is an enhancement to VPN technology to enable it to support the needs of new applications, particularly applications that are associated with 5G services. These applications require better isolation from both control and data plane's perspective and have more stringent performance requirements than can be provided with overlay VPNs. The characteristics of an enhanced VPN as perceived by its tenant needs to be comparable to those of a dedicated private network. This requires tight integration between the overlay VPN and the underlay network resources in a scalable manner. An enhanced VPN may form the underpinning of 5G network slicing, but will also be of use in its own right. This document describes the use of segment routing based mechanisms to provide the enhanced VPN service with dedicated network resources. The proposed mechanism is applicable to both SR with MPLS data plane and SR with IPv6 data plane (SRv6).

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Driven largely by needs arising from the 5G mobile network design, the concept of network slicing has gained traction [NGMN-NS-Concept] [TS23501][TS28530] [BBF-SD406]. Network slicing requires the transport network to support partitioning the network resources to provide the client with dedicated (private) networking, computing, and storage resources drawn from a shared pool. The slices may be seen as (and operated as) virtual networks.

Thus there is a need to create virtual networks with enhanced characteristics. The tenant of such a virtual network can require a degree of isolation and performance that previously could only be satisfied by dedicated networks. Additionally the tenant may ask for some level of control to their virtual network e.g. to customize the service paths in the network slice.

The enhanced VPN service (VPN+) as described in [I-D.ietf-teas-enhanced-vpn] is targeted at new applications which require better isolation from both control plane and data plane's perspective and have more stringent performance requirements than can be provided with existing overlay VPNs. An enhanced VPN may form the underpinning of network slicing, but will also be of use in its own right.

Although each VPN can be associated with a set of dedicated RSVP-TE [RFC3209] LSPs with bandwidth reservation to provide some guarantee to service performance, such mechanisms would introduce per-VPN per-path states into the network, which is known to have scalability issues [RFC5439] and has not been widely adopted in production networks.

Segment Routing (SR) [RFC8402] specifies a mechanism to steer packets through an ordered list of segments. It can achieve explicit source routing without introducing per-path state into the network. Like RSVP-TE, SR also supports source specification of the packet path. However, currently SR does not have the capability of reserving or identifying different network resources for different services or customers. Although the controller can have global view of network state and can provision different services onto different SR paths, in the data plane it still relies on traditional DiffServ QoS model [RFC2474] [RFC2475] to provide coarse-grained traffic differentiation in the network. While this may be sufficient for some traditional services, it cannot meet the requirement of the enhanced VPN service.

This document extends the SR paradigm by allocating different Segment Identifiers (SIDs) to represent the different subset of resources allocated on each network elements (links or nodes). The SIDs

associated with a particular group of network resources can be used to construct customized virtual networks for different services, the SID can also be used to steer the service traffic to be processed with the corresponding allocated resources. This mechanism can be used to provide the enhanced VPN service with dedicated network resources. The proposed mechanism is applicable to both SR with MPLS data plane and SR with IPv6 data plane (SRv6).

2. Requirements Notation

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 RFC 2119 [RFC2119] RFC8174 [RFC8174][when, and only when, they appear in all capitals, as shown here.

3. Segment Routing with Resource Awareness

In segment routing, several types of segments are defined to represent either topological elements or service instructions. A topological segment may be a node segment or an adjacency segment. Some other types of segments may be associated with specific service functions for service chaining purpose. However, so far none of the SR segments are associated with network resources for the QoS purpose.

In order to support the enhanced VPNs which require guaranteed performance and isolation from other services in the network, the overlay VPN needs to be integrated with part of the underlay networks. Some dedicated network resources need to be allocated to an enhanced VPN or a group of enhanced VPNs. When segment routing is used to provide enhanced VPNs, it is necessary to associate the segments with network resources. By extending the segment routing paradigm, different set of network resources can be allocated on network elements, and associated with different SIDs.

This section describes the possible mechanisms to bring resource-awareness into two SR data plane instantiations: SR-MPLS and SRv6.

3.1. SR-MPLS

3.1.1. Singe SID Identifying both Topology and Resource

In SR-MPLS [I-D.ietf-spring-segment-routing-mpls], Adjacency Segment (Adj-SID) is an IGP-segment attached to a unidirectional adjacency or a set of unidirectional adjacencies. Node segment is an IGP-Prefix segment that identifies a specific router (e.g., a loopback). These

two types of SIDs can be extended to represent both topological elements and the resources allocated on a particular network element.

On one particular network link, multiple adjacency segment identifiers (Adj-SIDs) can be allocated, each of which is associated with a subset of the link resource allocated, such as logical sub-interface, bandwidth, queues, etc. For one particular node, multiple node-SIDs can be allocated, each of which may be associated with a subset of resource allocated from the node, such as the processing resources. Per-segment resource allocation complies to the SR paradigm, which avoids introducing per-path state into the network.

Different groups of adj-SIDs and node-SIDs which represent different set of network resources can be used to build different virtual networks, which could be further used to provide different enhanced VPNs, so that the isolation and performance requirement of enhanced VPNs could be met. The adj-SIDs are used to steer traffic of different enhanced VPNs into different set of link resources. The node SIDs can be used to steer traffic of different enhanced VPNs into different node resources. The node SIDs can also be used to build loose SR paths for different enhanced VPNs. In this case, the node-SIDs are used by transit nodes to steer traffic into the local resources allocated for the corresponding enhanced VPN. Note in this case Penultimate Hop Popping (PHP) [RFC3031] MUST be disabled, as the node-SID is used to identify the SR virtual network and the corresponding network resources allocated to the enhanced VPN.

3.1.2. Dedicated SID Identifying Network Resource

Another option to bring resource-awareness into SR-MPLS data plane is to define a dedicated SID called "resource-SID" to identify the group of network resources allocated on a particular link or node. In SR label stack, the resource-SID MUST be encapsulated under the topological SIDs (adj-SID or node-SIDs) which identifies the network element it applies to.

Note that a network node can participate in multiple topologies. For each network topology it participates in, a dedicated node-SID is needed for topology-specific path computation and next hop resolution. Dedicated adj-SIDs could also be allocated for different network topologies.

In packet forwarding, the adj-SID and node-SID are used to determine the next-hop and the outbound interface in a particular virtual network, then the resource-SID is used to identify the fine granular forwarding plane resource to be used for the processing of the received packet.

The benefit of this approach is that it decouples the topology identification and resource identification. In some cases where multiple virtual networks share a same topology but map to different set of network resources, it is possible that the topology-specific processing (for example, SPF computation) could be shared, so that the scalability can be improved. The cost is it increases the depth of the MPLS label stack.

The resource-SID can be a global significant identifier, which represents the collection of network resources allocated in the whole network domain to a particular virtual network. In this case, the resource-SID SHOULD appear only once in the label stack, and it SHOULD be parsed by each transit node which performs per virtual network resource reservation. This resource-SID can be either a new type of SID, or it could be embedded in some existing MPLS labels. For example, some fields in the Entroy Label Indicator (ELI) / Entropy Label (EL) [RFC6790] may be used as the resource identifier, the details will be provided in a future version.

The resource-SID may be a local significant identifier, which only represents the network resource locally allocated on each network segment to a particular virtual network. In this case, it has to be added to the label stack for each hop which performs per-virtual network resource reservation. As this approach would increase the label stack depth significantly, this approach is NOT RECOMMENDED.

3.2. SRv6

An SRv6 Segment (SID) is a 128-bit value which consists of a locator (LOC) and a function (FUNCT), optionally it may also contain additional arguments (ARG) [I-D.filsfils-spring-srv6-network-programming]. The locator is used for routing towards a particular node, it needs to be parsed by all nodes in the network. The function and arguments are only parsed by the owner of the SRv6 SID to determine the local behavior on receipt of the SRv6 packet.

In order to build multiple virtual networks in an SRv6 network, each node SHOULD allocate a dedicated locator for each virtual network it participates in. In packet forwarding, the locator can be used to identify the virtual network the packet belongs to, so that a virtual network specific next-hop can be determined. In addition, the locator can also be used to identify the group of local network resources allocated to the virtual network. All the SRv6 functions associated with a particular virtual network MUST use the locator of that virtual network as the prefix to construct the SRv6 SID.

In some cases where multiple virtual networks share a same topology but maps to different set of network resources, it is possible that the topology-specific processing (for example, SPF computation) could be shared, so that the scalability can be improved. This requires to decouple the topology identification and resource identification in SRv6. The locator can still be used as the identifier of the topology, while another identifier is needed to identify the network resources allocated to a particular virtual network. There are some candidates for the resource identifier in the IPv6 [RFC8200] or SRv6 header [I-D.ietf-6man-segment-routing-header], such as the IPv6 Flow Label or the Hop-by-Hop Option. More details will be provided in a future version.

4. Control Plane

The architecture described in this document makes use of a centralized controller that collects the information about the network (configuration, state, routing databases, etc.) as well as the service information (traffic matrix, performance statistics, etc). The controller is also responsible for the centralized computation and optimization of the virtual networks used for enhanced VPNs. A distributed control plane is needed for the collection and distribution of the topology and state information of the virtual networks. Distributed routing computation for some services in the enhanced VPNs is also possible.

5. Procedures

This section describes the procedures of provisioning an enhanced VPN service based on segment routing with resource awareness.

According to the requirement of an enhanced VPN service, a centralized network controller calculates a subset of the underlay network topology to support this enhanced VPN. Within this topology, the network resources needed on each network element can also be determined. The network resources are allocated in a per-segment manner, and are associated with different node-SIDs and adj-SIDs. The group of the node-SIDs and adj-SIDs allocated for the enhanced VPN will be used by network nodes and the network controller to build a SR virtual topology, which is used as the logical underlay of the enhanced VPN service. The extensions to IGP protocol to distribute the SIDs and the associated resources allocated for a virtual network is specified in [I-D.dong-lsr-sr-enhanced-vpn].

Suppose that customer requests for an enhanced VPN service from the network operator. The fundamental requirement is that customer A's service does not experience interference from other services in the network, such as other customers' VPN services, or the non-VPN

services in the network. The detailed requirements can be described with characteristics such as the following:

- o Service topology: the service sites and the connectivity between them

- o Service bandwidth: the bandwidth requirement between service sites

- o Isolation: the level of isolation from other services in the network

- o Reliability: whether fast repair or end-to-end protection is needed or not.

- o Latency

- o Jitter

- o Visibility: the customer may want to have some form of visibility of the network deliversing the service.

5.1. Topology and Resource Computation

As described in section 4, a centralized network controller is responsible for the provisioning of enhanced VPNs. The controller needs to determine the information of network connectivity, network resources, network performance and other relevant network state of the underlay network. This is often done using either IGP [RFC5305] [RFC3630] [RFC7471] [RFC7810] or BGP-LS [RFC7752] [I-D.ietf-idr-te-pm-bgp].

Based on the network information collected from the underlay network, the controller computes the underlay topology (possibly using multiple algorithms) and knows the resources that are available and allocated. When a request is received from a tenant, the controller computes the subgraph of the underlay network, along with the resources to be allocated on each network element (e.g. links and nodes) in the topology to meet the tenant's requirements, whilst maintaining the needs of the existing tenants that are using the same network.

5.2. Network Resource and SID Allocation

According to the output of computation, the network controller instructs the network devices involved in the subgraph to allocate the required network resources for the enhanced VPN. 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, dedicated segment identifiers, e.g. node-SIDs and adj-SIDs are also allocated to represent the network resources allocated for the enhanced VPN on each network segment.

In the forwarding plane, there are multiple ways of allocating or reserving network resources to different enhanced VPNs. 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 of enhanced VPN can be found in [I-D.ietf-teas-enhanced-vpn]. The SR SIDs are used as a good abstraction of the various types of network resource reservation mechanisms in the forwarding plane.

```
    Node-SIDs:                          Node-SIDs:
      r:101                               r:102
      g:201   Adj-SIDs:                   g:202
      b:301      r:1001:1G    r:1001:1G   b:302
         +------+ g:2001:2G    g:2001:2G +------+
         |  A   | b:3001:1G    b:3001:1G |  B   |Adj-SIDs:
         |      +------------------------+      + r:1003:1G
Adj-SIDs +--+---+                        +--+---+\g:2003:2G
   r:1002:1G|                   r:1002:1G|    \
   g:2002:2G|                   g:2002:2G|     \ r:1001:1G
   b:3002:3G|                   b:3002:2G|      \g:2001:2G
            |                            |       \ +------+ Node-SIDs:
            |                            |        \+  E   |   r:105
            |                            |        /+      |   g:205
   r:1001:1G|                   r:1002:1G|       / +------+
   g:2001:2G|                   g:2002:2G|      /r:1002:1G
   b:3001:3G|                   b:3002:2G|     / g:2002:2G
         +--+---+                        +--+---+ /
         |      |                        |      |/r:1003:1G
         |  C   +------------------------+  D   + g:2003:2G
         +------+ r:1002:1G    r:1001:1G +------+
    Node-SIDs:   g:2002:1G    g:2001:1G   Node-SIDs:
      r:103      b:3002:2G    b:3001:2G     r:104
      g:203                                 g:204
      b:303                                 b:304
```

Figure 1. SIDs identify resources allocated to different virtual networks

Figure 1 shows a network fragment of enhanced VPN supported by SR. Note that the format of the SIDs in this figure are for illustration, both SR-MPLS and SRv6 can be utilized as the data plane. In this example, there are three virtual topologies created for enhanced VPNs red (r) , green (g) and blue (b). The red and green topologies consist of nodes A, B, C, D, and E with all their interconnecting links, whilst the blue topology only consists of nodes A, B, C and D

with all their interconnecting links. Each node allocates a dedicated adjacency SID for each link participating in a particular topology. Each node is also allocated with a dedicated node SID for each topology it participates in. The adj-SIDs are associated with the link resources (e.g. bandwidth) allocated to each topology, so that the adj-SIDs can be used to steer service of different enhanced VPNs into different set of reserved resources in the data plane. The node-SIDs can be associated with dedicated nodal resources allocated for each topology. In addition, the node-SIDs of different topologies can be used to build loose SR path within each virtual topology, and steer service of different enhanced VPNs into the different set of reserved resources in the data plane.

In Figure 1, the notation x:nnnn:y that in topology colour x, the adj-SID nnnn will steer the packet over that link which has a total bandwidth of y assigned to that topology. Thus the note r:1002:1G in link C->D says that the red topology over link C->D has a reserved bandwidth of 1Gb/s and will be used by packets arriving at node C with an adj-SID 1002 at the top of the label stack.

5.3. Construction of SR Virtual Networks

Each node MUST advertise its set of resources (allocated and available) and the associated SIDs both to the centralized controller and into the network. This can be achieve by many different means such as (non-exhaustive list) IGP extensions [I-D.dong-lsr-sr-enhanced-vpn], BGP-LS [RFC7752] with possible extensions, NETCONF/YANG [RFC6241] [RFC7950].

With the collected network resource and SIDs information, the controller and network nodes are able to construct the SR virtual topologies and forwarding entries using the node-SIDs and adj-SIDs allocated for each enhanced VPN. Unlike classic segment routing in which network resources are shared by all services and customers, the SR virtual networks are associated with dedicated resource allocated in the underlay, so that they can be used to meet the service requirement of enhanced VPN and provide the required isolation from other services in the same network.

Figure 2 shows the virtual SR topologies created from the underlay network in Figure 1.

```
      1001  1001                 2001  2001                 3001  3001
   101---------102            201---------202            301---------302
    |           | \1003        |           | \2003        |           |
1002|       1002|  \ 1001  2002|       2002|  \ 2001  3002|       3002|
    |           |  105         |           |  205         |           |
1001|       1002|  / 1002  2001|       2002|  / 2002  3001|       3002|
    |           | / 1003       |           | / 2003       |           |
   103---------104            203---------204            303---------304
      1002  1001                 1002  2001                 3002  3001
    Topology Red              Topology Green              Topology Blue
```

Figure 2. SR virtual topologies using different groups of SIDs

5.4. VPN Service to SR Virtual Network Mapping

The services of an enhanced VPN customer can be provisioned using the customized SR virtual network as the underlay. In this way, services of different enhanced VPNs will only use the network resources allocated and will not interfere with each other. For each enhanced VPN customer, the service paths can be customized for different services within the SR virtual topology, and the allocated network resources are shared by different services of the same enhanced VPN customer.

For example, to create a strict path along the path A-B-D-E in the red topology in Figure 2, the SR segment list in the service packet would be (1001, 1002, 1003). For the same strict path in green topology, the SR segment list would be (2001, 2002, 2003). In the case where we wish to construct a loose path A-D-E in the green topology, the service packet SHOULD be set with the SR segment list (201, 204, 205). At node A the packet is sent towards D via either node B or C using the link and node resources allocated for the green topology. At node D the packet is forwarded to E using the link and node resource allocated for the green topology. Similarly, a packet for the loose path A-D-E in the red topology would arrive at node A with the SID list (101, 104, 105).

5.5. Network Visibility to Customer

The tenants of enhanced VPNs 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 topology, or a set of computed paths with transit nodes, or simply the 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, the 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 network information.

6. Benefits of the Proposed Mechanism

The proposed mechanism provides several key characteristics:

- o Flexibility
- o Scalability
- o Resource isolation

In addition to isolation, the proposed mechanism allows resource sharing between different services of the same enhanced VPN customer. This gives the customer more flexibility and control in service planning and provisioning, the experience would be similar to using a dedicated private network. The performance of critical services flows in a particular enhanced VPN can be further ensured using the mechanisms defined in [DetNet].

The detailed comparison with other candidates technologies are given in the following subsections.

6.1. MPLS-TP

MPLS-TP could be enhanced to include the allocation of specific resources along the path to a specific LSP. This would require that the SDN system set up and maintain every resource at every path for every customer, and map this to the LSP in the data plane, hence at every hop unique LSP label is needed for each path. Whilst this would be a way to produce a proof of concept for network slicing of an MPLS underlay, delegation would be difficult, resulting in a high overhead and a system needing too much administration. This leads to scaling concerns. The number of labels needed at any node would be the total number of services passing through that node. Experience with early pseudowire designs shows that this can lead to scaling issues.

6.2. RSVP-TE

RSVP-TE has the same scaling concern as MPLS-TP in terms of the number of LSPs that need to be maintained being equal to the number of services passing through any given node. It also has the two RSVP disadvantages that basic SR seeks to address:

- The use of RSVP for path establishment in addition to the routing protocol used to discover the topology and the network resources.

- The overhead of the soft-state maintenance associated with RSVP. The impact of this overhead would be exacerbated by the increased number of end to end paths requiring state maintenance.

6.3. Basic SR

Compared to RSVP, SR reduces the number of control protocols to only the routing protocol. It also attempts to minimize the core state by pushing state into the packet, although in some cases the binding SIDs are required to overcome the limitations in the ability of some nodes to push large label stacks. Moreover, currently SR does not support resource allocation or identification below the level of link, and none at node level. This restricts the extent to which some particular tenant traffic can be isolated from other traffic in the network.

6.4. SR with Resource Awareness

The approach described in this document seeks to achieve a compromise between the state limitations of traditional TE systems and the lack of resource awareness in basic SR.

By segmenting the path and allocating network resources to each element of the virtual network topologies, the operator can choose the granularity of resource to path binding within a virtual topology. In network segments where resource is scarce such that the service requirement may not always be met, the SR approach can allocate specific resources to a particular high priority service. By contrast, in other parts of the network where resource is plentiful, the resource may be shared by a number of services. The decision to do this is in the hands of the operator. Because of the segmented nature of the path, resource aggregation is possible in a way that is more difficult with RSVP-TE and MPLS-TP due to the use of dedicated label to identify each end-to-end path.

7. Service Assurance

In order to provide service assurance it is necessary to instrument the network at multiple levels. The network operator needs to ascertain that the underlay is operating correctly. A tenant needs to ascertain that their services are correctly operating. 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. 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.

8. IANA Considerations

This document makes no request of IANA.

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

9. Security Considerations

The normal security considerations of VPNs are applicable and it is assumed that industry best practise is applied to an enhanced VPN.

The security considerations of segment routing are applicable and it is assumed that these are applied to an enhanced VPN that uses SR.

Some applications of enhanced VPNs are sensitive to packet latency; the enhanced VPNs provisioned to carry their traffic have 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 critical services this type of attack can be prevented. However case 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 MUST NOT be exposed to third parties, to prevent attacks aimed at exploiting a shared resource.

10. Acknowledgements

The authors would like to thank Mach Chen, Zhenbin Li, Stefano Previdi, Charlie Perkins and Bruno Decraene for the 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-18 (work in progress), December 2018.

[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

[BBF-SD406]
"BBF SD-406: End-to-End Network Slicing", 2016, <https://wiki.broadband-forum.org/display/BBF/SD-406+End-to-End+Network+Slicing>.

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

[I-D.dong-lsr-sr-enhanced-vpn]
Dong, J. and S. Bryant, "IGP Extensions for Segment Routing based Enhanced VPN", draft-dong-lsr-sr-enhanced-vpn-01 (work in progress), October 2018.

[I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-07 (work in progress), February 2019.

[I-D.ietf-6man-segment-routing-header]
Filsfils, C., Previdi, S., Leddy, J., Matsushima, S., and d. daniel.voyer@bell.ca, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-16 (work in progress), February 2019.

[I-D.ietf-idr-te-pm-bgp]
Ginsberg, L., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions", draft-ietf-idr-te-pm-bgp-18 (work in progress), December 2018.

[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-01 (work in progress), February 2019.

[NGMN-NS-Concept]
"NGMN NS Concept", 2016, <https://www.ngmn.org/fileadmin/user_upload/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf>.

[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>.

[TS23501] "3GPP TS23.501", 2016, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144>.

[TS28530] "3GPP TS28.530", 2016, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3273>.

Authors' Addresses

Jie Dong
Huawei Technologies

Email: jie.dong@huawei.com

Stewart Bryant
Huawei Technologies

Email: stewart.bryant@gmail.com

Zhenqiang Li
China Mobile

Email: li_zhenqiang@hotmail.com

Takuya Miyasaka
KDDI Corporation

Email: ta-miyasaka@kddi.com