Authors: J. Dong              S. Bryant
         Huawei Technologies   University of Surrey
         T. Miyasaka       Y. Zhu          F. Qin
         KDDI Corporation   China Telecom   China Mobile
         Z. Li          F. Clad
         China Mobile   Cisco Systems

# Segment Routing based Virtual Transport Network (VTN) for Enhanced VPN

## Abstract

Segment Routing (SR) leverages the source routing paradigm. A node
steers a packet through an ordered list of instructions, called
"segments". A segment can represent topological or service based
instructions. A segment can further be associated with a set of
network resources used for executing the instruction. Such a segment
is called resource-aware segment.

A Virtual Transport Network (VTN) is a virtual underlay network
which has a network topology and a set of network resources
allocated from the physical underlay network.

Resource-aware Segment Identifiers (SIDs) may be used to build SR
paths with a set of reserved network resources. In addition, a group
of resource-aware SIDs may be used to build SR based virtual
underlay networks, which have customized network topology and
resource attributes required by one or a group of customers and/or
services. Such virtual networks are the SR instantiations of VTNs.

This document describes a suggested way to use resource-aware SIDs
to build SR based VTNs.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the
provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering
Task Force (IETF). Note that other groups may also distribute
working documents as Internet-Drafts. The list of current Internet-
Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six
months and may be updated, replaced, or obsoleted by other documents

at any time. It is inappropriate to use Internet-Drafts as reference
material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 April 2023.

**Copyright Notice**

**Table of Contents**

1.  **Introduction**

Segment Routing (SR) [RFC8402] specifies a mechanism to steer
packets through an ordered list of segments. A segment is referred
to by its Segment Identifier (SID). With SR, explicit source routing

can be achieved without introducing per-path state into the network. [I-D.ietf-spring-resource-aware-segments] extends SR by associating SIDs with network resource attributes (e.g., bandwidth, processing or storage resources). 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. Multiple resource-aware SIDs may be allocated on a network segment, each associated with a set of network resources assigned to meet the requirements of one or a group of customers and/or services.

A Virtual Transport Network (VTN) is a virtual underlay network which has a network topology and a set of network resources allocated from the physical underlay network.

Once allocated, Resource-aware SIDs can be used to build SR paths with a set of reserved network resources. In addition, a group of resource-aware SIDs may be used to build SR based virtual networks, which have customized network topology and resource attributes required by one or a group of customers and/or services. Such virtual networks are the SR instantiations of VTNs as defined in [I-D.ietf-teas-enhanced-vpn], and can be used to enable the enhanced VPN (VPN+) services.

This document describes a suggested way to use resource-aware SIDs to build SR based VTNs. Although the procedure is illustrated using SR-MPLS, this mechanism is applicable to both SR over MPLS data plane (SR-MPLS) [RFC8660] and SR over IPv6 data plane (SRv6) [RFC8986].

## 2.  Resource-Aware SIDs for VTN

As defined in [I-D.ietf-teas-enhanced-vpn], a VTN is a virtual underlay network which has a specific network topology and a subset of network resources allocated from the physical network.

When SR is used as the data plane of VTNs in the network, it is necessary to compute and instantiate the SR paths with the topology and/or algorithm constraints of the VTN, and steer the traffic to only use the set of network resources allocated to the VTN.

Based on the resource-aware segments defined in [I-D.ietf-spring-resource-aware-segments], a group of resource-aware SIDs can be allocated to represent the network segments of a VTN. These resource-aware SIDs are associated with the group of network resources allocated to the VTN on network nodes and links which participate in the VTN. These resource-aware SIDs can also identify the network topological or functional instructions associated with the VTN.

The resource-aware SIDs may be allocated either by a centralized
network controller or by network nodes. The control plane mechanisms
for advertising the resource-aware SIDs for VTNs can be based on
[RFC4915], [RFC5120] and [I-D.ietf-lsr-flex-algo] with necessary
extensions. This is further described in section 3.3.

## 2.1.  SR-MPLS based VTN

This section describes a mechanism of allocating resource-aware SIDs
to SR-MPLS based VTNs.

For an IGP link, multiple adj-SIDs are allocated, each of which is
associated with a VTN that the link participates in, and represents
those link resources that are allocated to the VTN. For an IGP node,
multiple prefix-SIDs are allocated, each of which is associated with
a VTN which the node participates in, and identifies the sets of
network resources allocated to the VTN on network nodes which
participate in the VTN. These set of resources will be used to
process packets which have the resource-aware SIDs as the active
segment.

In the case of multi-domain VTNs, on an inter-domain link, multiple
BGP peering SIDs [RFC9086] are allocated, each of which is
associated with a VTN which spans multiple domains, and represents a
subset of resources allocated on the inter-domain link.

## 2.2.  SRv6 based VTN

This section describes a mechanism of allocating resource-aware SRv6
Locators and SIDs to SRv6 based VTNs.

An SRv6 Locator is allocated on each network node for each VTN in
which the node participates. This Locator (the VTN-specific Locator)
identifies the set of network resources allocated to the VTN on the
network nodes which participate in the VTN. The SRv6 SIDs associated
with a VTN are allocated from the SID space using the VTN-specific
Locator as the prefix. These SRv6 SIDs can be used to represent VTN-
specific SRv6 functions, and can identify the set of resources used
by network nodes to process packets.

## 2.3.  VTN Identification

In a simple case, each VTN can be mapped to a unique topology or
algorithm. Then the VTNs can be distinguished by the topology ID or
algorithm ID in the control plane, and the resource-aware SIDs
associated with a VTN can be identified using the <topology,
algorithm> tuple as described in [RFC8402]. In this case, the number
of VTNs supported in a network relies on the number of topologies or
algorithms supported.

In a more complicated case, multiple VTNs may be mapped to the same
<topology, algorithm> tuple, while each is allocated a separate set
of network resources. Then a new VTN identifier (VTN-ID) in the
control plane is needed to identify the VTN. The resource-aware SIDs
associated with different VTNs can be distinguished using VTN-IDs in
the control plane.

In the data plane, the resource-aware SIDs are used to identify the
VTN, and are also used to determine the forwarding instructions and
the set of network resources used for the packet processing action.

## 2.4.  Scalability Considerations

Since multiple VTNs can be created in a network, and each VTN is
allocated a group of resource-aware SIDs, the mechanism of SR based
VTNs increases the number of SIDs and SRv6 Locators needed in a
network. There may be some concern, especially about the SR-MPLS
prefix-SIDs, which are allocated from the Segment Routing Global
Block (SRGB), that the SRGB will be used up. The amount of network
state will also increase accordingly. However, based on the SR
paradigm, resource-aware SIDs and the associated network state are
allocated and maintained per VTN, thus per-path network state is
avoided in the SR network. In the control plane, the amount of
information to be distributed for different VTNs may also become a
concern for the IGP protocols. The scalability of resource-aware SID
based VTNs are further analysed in [I-D.ietf-teas-nrp-scalability].

## 3.  Procedures

This section describes possible procedures for creating SR based
VTNs and the corresponding forwarding tables and entries. The
approaches described in this section are not normative, but
illustrate how the VTN-specific Locator and VTN ID could be used to
build and operate VTNs in SR networks. Although it is illustrated
using SR-MPLS, this mechanism is applicable to both SR-MPLS and
SRv6.

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

According to the customer's requirements, a centralized network
controller calculates a subset of the underlay network topology to
support the service. With this topology, the set of network
resources required on each network element is also determined. The
subset of network topology and network resources are the two major

characteristics of a VTN. Depending on the service requirements, the network topology and network resource of this VTN can be dedicated for an individual customer or service, or can be shared by a group of customers and/or services.

Based on the mechanisms described in section 2, a group of resource-aware SIDs can be allocated for the VTN. With SR-MPLS, it is a group of prefix-SIDs and adj-SIDs which are allocated to identify the network nodes and links in the VTN, and also to identify the set of network resources allocated on these network nodes and links for the VTN. As the resource-aware SIDs can be allocated either by a centralized network controller or by the network nodes, control plane protocols such as IGP (e.g., IS-IS or OSPF) and BGP-LS can be used to distribute the SIDs and the associated resource and topology information of a VTN to other nodes in the same VTN and also to the controller, so that both the network nodes and the controller can generate the VTN-specific forwarding table or forwarding entries based on the resource-aware SIDs of the VTN. The detailed control plane mechanisms and possible extensions are described in the accompanying documents [I-D.ietf-lsr-isis-sr-vtn-mt] [I-D.ietf-idr-bgpls-sr-vtn-mt] [I-D.zhu-lsr-isis-sr-vtn-flexalgo] [I-D.zhu-idr-bgpls-sr-vtn-flexalgo] [I-D.dong-lsr-sr-enhanced-vpn] [I-D.dong-idr-bgpls-sr-enhanced-vpn] and are out of the scope of this document.

## 3.1. VTN Topology and Resource Planning

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

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

## 3.2. VTN Network Resource and SID Allocation

According to the result of VTN planning, the network controller
instructs the set of network nodes involved to join a specific VTN
and allocate the required set of network resources for the VTN. This
may be done with Netconf/YANG [RFC6241] [RFC7950] or with any other
control or management plane mechanism with necessary extensions.
Thus, the controller not only allocates the resources to the newly
computed VTN, but also keeps track of the remaining available
resources in order to cope with subsequent VTN requests.

On each network node involved in the VTN, a set of network resources
(e.g., link bandwidth) is allocated to the VTN. Such set of network
resources can be dedicated for the processing of traffic in that
VTN, and cannot be used by traffic in other VTNs. Note it is also
possible that a group of VTNs may share a set of network resources
on some network segments. A group of resource-aware SIDs, such as
prefix-SIDs and adj-SIDs are allocated to identify both the network
segments and the set of resources allocated on the network segments
for the VTN. Such group of resource-aware SIDs, e.g., prefix-SIDs
and adj-SIDs are used as the data plane identifiers of the nodes and
links in the VTN.

In the underlying forwarding plane, there can be multiple ways of
allocating a subset of network resources to a VTN. The candidate
data plane technologies to support resource partitioning or
reservation can be found in [I-D.ietf-teas-enhanced-vpn]. The
resource-aware SIDs are considered as abstract data plane
identifiers in the network layer, which can work with various
network resource partitioning or reservation mechanisms in the
underlying forwarding plane.

```
   Prefix-SIDs:                        Prefix-SIDs:
     r:101                                r:102
     g:201                                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  |
        |     +------------------------+     + r:1003:1G
        +--+--+                        +--+--+\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
           |                             |        \ +-----+Prefix-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 +-----+
   Prefix-SIDs:  g:2002:1G    g:2001:1G  Prefix-SIDs:
     r:103       b:3002:2G    b:3001:2G     r:104
     g:203                                  g:204
     b:303                                  b:304
```
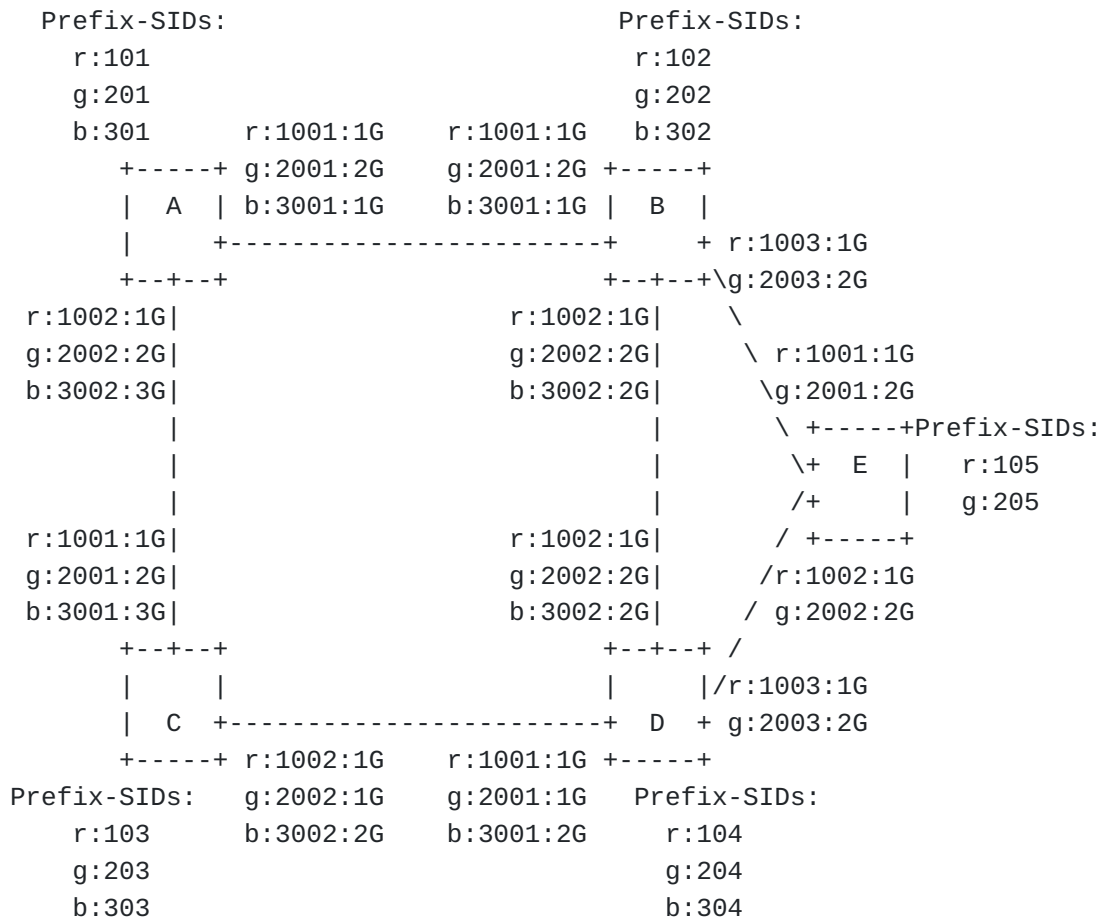
    Figure 1. SID and resource allocation for multiple VTNs

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

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

resource-aware adj-SID can be associated with a set of link
resources (e.g., bandwidth) allocated to a specific VTN, so that
different adj-SIDs can be used to steer service traffic into
different set of link resources in packet forwarding. A resource-
aware prefix-SIDs in a VTN can be associated with the set of network
resources allocated to this VTN on each involved network node and
link. Thus, the prefix-SIDs can be used to build loose SR path
within a VTN, and can be used by the transit nodes to steer traffic
into the set of local network resources allocated to the VTN.

### 3.3.  Construction of SR based VTNs

The network controller needs to obtain the information about all the
VTNs in the network it oversees, including the resource-aware SIDs
and their associated network resources and topology information.
Based on this information, the controller can have a global view of
the VTN topologies, network resources and the associated SIDs, so as
to perform VTN-specific explicit path computation, taking both the
topology and resource constraints of the VTNs into consideration,
and use the resource-aware SIDs to build the SID list for the
explicit path. The controller may also compute the shortest paths in
the VTN based on the resource-aware prefix-SIDs.

The network nodes also need to obtain the information about the VTNs
they participate in, including the resource-aware SIDs and their
associated network resources and topology information. Based on the
collected information, the network nodes which are the headend of a
path can perform VTN-specific path computation, and build the SID
list using the collected resource-aware adj-SIDs and prefix-SIDs.
The network nodes also need to generate the forwarding entries for
the resource-aware prefix-SIDs in each VTN they participates in, and
associate these forwarding entries with the set of local network
resources (e.g., bandwidth on the outgoing interface) allocated to
the corresponding VTN.

Thus, after receiving the network controller's instruction about
network resource and SID allocation, each network node needs to
advertise the identifier of the VTNs it participates in, the group
of resource-aware SIDs allocated to each VTN, and the resource
attributes (e.g., bandwidth) associated with the resource-aware SIDs
in the network. Each resource-aware adj-SID is advertised with the
set of associated link resources, and each resource-aware prefix-SID
is advertised with the identifier of the associated VTN, as all the
prefix-SIDs in a VTN are associated with the same set of network
resources allocated to the VTN. Note that, as described in section
2.3, the VTNs can be identified in the control plane either using
existing identifiers, such as the MT-ID or Flex-Algo ID, or using a
newly defined VTN ID.

The IGP mechanisms which reuse the existing IDs such as Multi-Topology [RFC5120] or Flex-Algo [I-D.ietf-lsr-flex-algo] as the identifier of VTNs, and distribute the resource-aware SIDs and the associated topology and resource information may be based on the mechanisms described in [I-D.ietf-lsr-isis-sr-vtn-mt] and [I-D.zhu-lsr-isis-sr-vtn-flexalgo] respectively. The corresponding BGP-LS mechanisms which can be used to distribute both the intra-domain VTN information and the inter-domain VTN-specific link information to the controller may be based on the mechanisms described in [I-D.ietf-idr-bgpls-sr-vtn-mt] and [I-D.zhu-idr-bgpls-sr-vtn-flexalgo] respectively. Note that with these mechanisms, the number of VTNs supported relies on the number of topologies or algorithms supported.

The IGP mechanisms described in [I-D.dong-lsr-sr-enhanced-vpn] introduce a new VTN-ID in the control plane, so that multiple VTNs can be mapped to the same <topology, algorithm> tuple, while each VTN can have different resource attributes. This provides a mechanism which allows flexible combination of network topology and network resources attributes to build a large number of VTNs with a relatively small number of topologies or algorithms. The corresponding BGP-LS mechanisms which may be used to distribute the intra-domain VTN information and the inter-domain VTN-specific link information to the controller are described in [I-D.dong-idr-bgpls-sr-enhanced-vpn].

Figure 2 shows the three SR based VTNs created in the 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                  2002  2001                  3002  3001
     VTN Red                     VTN Green                    VTN Blue
```

Figure 2. SR based VTNs with different groups of SIDs

For each SR based VTN, SR paths are computed within the VTN, taking the VTN topology and resources as constraints. The SR path can be an explicit path instantiated using SR policy [RFC9256], in which the SID-list is built only with the SIDs allocated to the VTN. The SR path can also be an IGP computed path associated with a prefix-SID or SRv6 End SID allocated by a node for the VTN, the IGP path computation is also based on the topology and/or algorithm

constraints of the VTN. Different SR paths in the same VTN may use
shared network resources when they use the same resource-aware SIDs
allocated to the VTN, while SR paths in different VTNs use different
set of network resources even when they traverse the same network
links or nodes. These VTN-specific SR paths need to be installed in
the corresponding forwarding tables.

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

## 3.4.  Mapping Services to SR based VTN

Network services can be provisioned using SR based VTNs as the
virtual underlay networks. For example, different services may be
provisioned in different SR based VTNs, each of which would use the
network resources allocated to the VTN, so that their data traffic
will not interfere with each other. In another case, a group of
services which have similar characteristics and requirements may be
provisioned in the same VTN, in this case the network resources
allocated to the VTN are only shared among this group of services,
but will not be shared with other services in the network. The
steering of service traffic to SR based VTNs can be based on either
local policy or, for example, the mechanisms as defined in
[RFC9256].

## 3.5.  VTN Visibility to Customers

VTNs can be used by network operators to organize and split their
network infrastructure into different virtual underlay networks for
different customers or services. Some customers may also request
different granularity of visibility into the VTN which is used to
deliver the service. Depending on the requirement and the network
operator's policy, VTNs may be exposed to the customer either as a
virtual network with both the edge nodes and the intermediate nodes,
as a set of paths with some of the transit nodes, or simply as a set
of virtual connections between the endpoints without any transit
node information. The visibility may be delivered through different

mechanisms, such as IGPs (e.g., IS-IS, OSPF), BGP-LS or Netconf/
YANG. On the other hand, a network operator may want to restrict the
visibility of the underlay network information it delivers to the
customer by either hiding the transit nodes between sites (and only
delivering information about the endpoint connectivity), or by
hiding some of the transit nodes (summarizing the path into fewer
nodes). The information about VTNs which are not used by the
customer should also be filtered. Mechanisms such as BGP-LS allow
the flexibility of the advertisement of aggregated virtual network
information and configurable filtering policies.

## 4.  Characteristics of SR based VTNs

The mechanism described in this document provides several key
characteristics:

  *Customization: Different customized VTNs can be created in a
   shared network to meet different customers' connectivity and
   service requirement. The customers are only aware of the topology
   and attributes of their own VTNs, and provision services on the
   VTN instead of the physical network. This provides a practical
   mechanism to support network slicing [I-D.ietf-teas-ietf-network-
   slices].

  *Resource Isolation: The computation and instantiation of SR paths
   in one VTN can be independent from other VTNs or other services
   in the network. In addition, a VTN can be associated with a set
   of dedicated network resources, which can avoid resource
   competition and performance interference from other VTNs or other
   services in the network. This mechanism also allows resource
   sharing between different service flows of the same customer, or
   between a group of services which are provisioned in the same
   VTN. This gives the operators and the customers the flexibility
   in network planning and service provisioning. In a VTN, the
   performance of critical services can be further ensured using
   other mechanisms, e.g., those as defined in [DetNet].

  *Scalability: The introduction of resource aware SIDs for
   different VTNs would increase the amount of SIDs and state in the
   network. While the increased network state is considered an
   inevitable price in meeting the requirements of some customers or
   services, the SR based VTN mechanism seeks to achieve a balance
   between the state limitations of traditional end-to-end TE
   mechanism and the lack of resource awareness in classic segment
   routing. Following the segment routing paradigm, network
   resources are allocated on network segments in a per VTN manner
   and represented as SIDs, this ensures that there is no per-path
   state introduced in the network. In addition, operators can
   choose the granularity of resource allocation on different

network segments. In network segments where resource is scarce such that the service requirement may not always be met, this approach can be used to allocate a set of resources to a VTN which contains such network segment to avoid possible competition. By contrast, in other segment of the network where resource is considered plentiful, the resource may be shared between a number of VTNs. The decision to do this is in the hands of the operator.

## 5.  Service Assurance of VTNs

In order to provide assurance for services provisioned in the SR based VTNs, it is necessary to instrument the network at multiple levels, e.g., in both the underlay network level and the VTN level. The operator or the customer may also monitor and measure the performance of the services carried by the VTN. In principle these can be achieved using existing or in development techniques in IETF, such as network telemetry [RFC9232]. The detailed mechanisms are out of the scope of this document.

In case of failure or service performance degradation in a VTN, it is necessary that some recovery mechanisms, e.g., local protection or end-to-end protection mechanism is used to switch the traffic to another path in the same VTN which could meet the service performance requirement. Care must be taken that the service or path recovery mechanism in one VTN does not impact other VTNs in the same network.

## 6.  IANA Considerations

This document makes no request of IANA.

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

## 7.  Security Considerations

The security considerations of segment routing [RFC8402] [RFC8754] and resource-aware SIDs [I-D.ietf-spring-resource-aware-segments] are applicable to this document.

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

and when the network needs to be reconfigured as part of ongoing maintenance or in response to a failure.

Considering the scalability of the SR VTN mechanism, the system may be destabilised by an attack or accident) that causes a large number of VTNs to be configured. This can be mitigated by placing thresholds (for alarms or cut-off) in the configuration process.

Traffic within a network may marked as belonging to a specific VTN and this makes it possible to carry out targetted attacks on traffic and to deduce customer-sensitive traffic patterns.

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

## 8. Contributors

Zhenbin Li
Email: lizhenbin@huawei.com

Zhibo Hu
Email: huzhibo@huawei.com

## 9. Acknowledgements

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

## 10. References

## 10.1. Normative References

**[I-D.ietf-spring-resource-aware-segments]**
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., Li, Z., and F. Clad, "Introducing Resource Awareness to SR Segments", Work in Progress, Internet-Draft, draft-ietf-spring-resource-aware-segments-06, 11 October 2022, <https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-spring-resource-aware-segments/>.

**[I-D.ietf-teas-enhanced-vpn]** Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Network (VPN+)", Work in Progress, Internet-Draft, draft-ietf-teas-enhanced-vpn-11, 19 September 2022, <https://www.ietf.org/archive/id/draft-ietf-teas-enhanced-vpn-11.txt>.

**[RFC8402]**
        Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

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

## 10.2.  Informative References

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

**[I-D.dong-idr-bgpls-sr-enhanced-vpn]** Dong, J., Hu, Z., Li, Z., Tang, X., and R. Pang, "BGP-LS Extensions for Scalable Segment Routing based Enhanced VPN", Work in Progress, Internet-Draft, draft-dong-idr-bgpls-sr-enhanced-vpn-04, 4 March 2022, <https://www.ietf.org/archive/id/draft-dong-idr-bgpls-sr-enhanced-vpn-04.txt>.

**[I-D.dong-lsr-sr-enhanced-vpn]** Dong, J., Hu, Z., Li, Z., Tang, X., Pang, R., and S. Bryant, "IGP Extensions for Scalable Segment Routing based Enhanced VPN", Work in Progress, Internet-Draft, draft-dong-lsr-sr-enhanced-vpn-08, 11 July 2022, <https://www.ietf.org/archive/id/draft-dong-lsr-sr-enhanced-vpn-08.txt>.

**[I-D.ietf-idr-bgpls-sr-vtn-mt]** Xie, C., Li, C., Dong, J., and Z. Li, "BGP-LS with Multi-topology for Segment Routing based Virtual Transport Networks", Work in Progress, Internet-Draft, draft-ietf-idr-bgpls-sr-vtn-mt-01, 12 September 2022, <https://www.ietf.org/archive/id/draft-ietf-idr-bgpls-sr-vtn-mt-01.txt>.

**[I-D.ietf-lsr-flex-algo]** Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-25, 6 October 2022, <https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-25.txt>.

**[I-D.ietf-lsr-isis-sr-vtn-mt]** Xie, C., Ma, C., Dong, J., and Z. Li, "Using IS-IS Multi-Topology (MT) for Segment Routing based Virtual Transport Network", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-sr-vtn-mt-03, 10 July 2022, <https://www.ietf.org/archive/id/draft-ietf-lsr-isis-sr-vtn-mt-03.txt>.

**[I-D.ietf-teas-ietf-network-slices]**
        Farrel, A., Drake, J., Rokui, R., Homma, S., Makhijani,
        K., Contreras, L. M., and J. Tantsura, "Framework for
        IETF Network Slices", Work in Progress, Internet-Draft,
        draft-ietf-teas-ietf-network-slices-14, 3 August 2022,
        <https://www.ietf.org/archive/id/draft-ietf-teas-ietf-
        network-slices-14.txt>.

**[I-D.ietf-teas-nrp-scalability]**
        Dong, J., Li, Z., Gong, L., Yang, G., Guichard, J. N.,
        Mishra, G., Qin, F., Saad, T., and V. P. Beeram,
        "Scalability Considerations for Network Resource
        Partition", Work in Progress, Internet-Draft, draft-ietf-
        teas-nrp-scalability-00, 11 July 2022, <https://
        www.ietf.org/archive/id/draft-ietf-teas-nrp-
        scalability-00.txt>.

**[I-D.zhu-idr-bgpls-sr-vtn-flexalgo]** Zhu, Y., Dong, J., and Z. Hu,
        "BGP-LS with Flex-Algo for Segment Routing based Virtual
        Transport Networks", Work in Progress, Internet-Draft,
        draft-zhu-idr-bgpls-sr-vtn-flexalgo-01, 22 February 2021,
        <https://www.ietf.org/archive/id/draft-zhu-idr-bgpls-sr-
        vtn-flexalgo-01.txt>.

**[I-D.zhu-lsr-isis-sr-vtn-flexalgo]** Zhu, Y., Dong, J., and Z. Hu,
        "Using Flex-Algo for Segment Routing (SR) based Virtual
        Transport Network (VTN)", Work in Progress, Internet-
        Draft, draft-zhu-lsr-isis-sr-vtn-flexalgo-05, 11 July
        2022, <https://www.ietf.org/archive/id/draft-zhu-lsr-
        isis-sr-vtn-flexalgo-05.txt>.

**[RFC3630]**  Katz, D., Kompella, K., and D. Yeung, "Traffic
        Engineering (TE) Extensions to OSPF Version 2", RFC 3630,
        DOI 10.17487/RFC3630, September 2003, <https://www.rfc-
        editor.org/info/rfc3630>.

**[RFC4915]**  Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P.
        Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF",
        RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://
        www.rfc-editor.org/info/rfc4915>.

**[RFC5120]**  Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi
        Topology (MT) Routing in Intermediate System to
        Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/

RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[RFC5305]  Li, T. and H. Smit, "IS-IS Extensions for Traffic
           Engineering", RFC 5305, DOI 10.17487/RFC5305, October
           2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC6241]  Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J.,
           Ed., and A. Bierman, Ed., "Network Configuration Protocol
           (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011,
           <https://www.rfc-editor.org/info/rfc6241>.

[RFC7471]  Giacalone, S., Ward, D., Drake, J., Atlas, A., and S.
           Previdi, "OSPF Traffic Engineering (TE) Metric
           Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015,
           <https://www.rfc-editor.org/info/rfc7471>.

[RFC7752]  Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A.,
           and S. Ray, "North-Bound Distribution of Link-State and
           Traffic Engineering (TE) Information Using BGP", RFC
           7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[RFC7950]  Bjorklund, M., Ed., "The YANG 1.1 Data Modeling
           Language", RFC 7950, DOI 10.17487/RFC7950, August 2016,
           <https://www.rfc-editor.org/info/rfc7950>.

[RFC8570]  Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward,
           D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE)
           Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March
           2019, <https://www.rfc-editor.org/info/rfc8570>.

[RFC8571]  Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J.,
           and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement
           of IGP Traffic Engineering Performance Metric
           Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019,
           <https://www.rfc-editor.org/info/rfc8571>.

[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>.

     [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>.

     [RFC9232]    Song, H., Qin, F., Martinez-Julia, P., Ciavaglia, L.,
                  and A. Wang, "Network Telemetry Framework", RFC 9232, DOI
                  10.17487/RFC9232, May 2022, <https://www.rfc-editor.org/
                  info/rfc9232>.

     [RFC9256]    Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov,
                  A., and P. Mattes, "Segment Routing Policy Architecture",
                  RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://
                  www.rfc-editor.org/info/rfc9256>.

Authors' Addresses

   Jie Dong
   Huawei Technologies

   Email: jie.dong@huawei.com

   Stewart Bryant
   University of Surrey

   Email: stewart.bryant@gmail.com

   Takuya Miyasaka
   KDDI Corporation

   Email: ta-miyasaka@kddi.com

   Yongqing Zhu
   China Telecom

   Email: zhuyq8@chinatelecom.cn

   Fengwei Qin
   China Mobile

   Email: qinfengwei@chinamobile.com

   Zhenqiang Li
   China Mobile

   Email: li_zhenqiang@hotmail.com

Francois Clad
Cisco Systems

Email: fclad@cisco.com