

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 26, 2021

W. Britto
R. Shetty
C. Barth
Juniper Networks
B. Wen
Comcast
S. Peng
R. Chen
ZTE Corporation
February 22, 2021

IGP Extensions for Support of Slice Aggregate Aware Traffic Engineering
[draft-bestbar-lsr-slice-aware-te-00](#)

Abstract

A slice aggregate is a collection of packets that match a slice policy selection criteria and are given the same forwarding treatment. Slice Aggregate aware Traffic Engineering (SA-TE) is a mechanism that facilitates Traffic Engineering (TE) path selection to take into account the available network resources associated with a specific slice aggregate. This document specifies the Interior Gateway Protocol (IGP) extensions for support of SA-TE. This includes the generalization of the semantics of a number of IGP extensions already defined for existing MPLS Traffic Engineering in [RFC3630], [RFC4124], [RFC5305] and additional IGP extensions beyond those.

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 August 26, 2021.

Copyright Notice

Copyright (c) 2021 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	3
2.	Requirements Language	4
3.	Terminology	4
4.	Resource Allocation	4
4.1.	Resource Sharing and Resource Overbooking	5
4.2.	SA-TE Resource Allocation Model	5
5.	Next-hop Filtering Capability	5
6.	IS-IS Extensions for SA-TE	7
6.1.	Network Slicing GENINFO TLV	7
6.2.	Slice Aggregate Aware Traffic Engineering Capabilities APPSub-TLV	7
6.3.	Slice Aggregate Aware Traffic Engineering Link APPSub-TLV	8
6.4.	Slice Aggregate Aware Traffic Engineering Link Attributes Sub-sub-TLV	10
6.4.1.	Slice Aggregate Aware Traffic Engineering Unreserved Bandwidth Sub-sub-sub-TLV	10
6.4.2.	Slice Aggregate Aware Traffic Engineering Residual Bandwidth Sub-sub-sub-TLV	11
6.4.3.	Slice Aggregate Aware Traffic Engineering Available Bandwidth Sub-sub-sub-TLV	11
6.4.4.	Slice Aggregate Aware Traffic Engineering Utilized Bandwidth Sub-sub-sub-TLV	12
7.	OSPF Extensions for SA-TE	12
8.	Functional Example of Unreserved-Bandwidth Advertisements . .	12
8.1.	Formula to Compute the Unreserved Bandwidth	12
8.2.	Examples	13
8.2.1.	Example 1 Results	15
8.2.2.	Example 2 Results	16
8.2.3.	Example 3 Results	17
9.	Scaling Considerations	18
9.1.	Link-State Update Frequency	18

10.	IANA Considerations	18
10.1.	Network Slicing GENINFO TLV	18
10.2.	Network Slicing GENINFO APPsub-TLVs	18
10.3.	SA-TE Link Sub-sub-TLVs	19
10.4.	SA-TE Link Attributes Sub-sub-sub-TLVs	19
11.	Security Considerations	20
12.	References	20
12.1.	Normative References	20
12.2.	Informative References	22
Appendix A.	Contributors	22
Appendix B.	Acknowledgements	23
	Authors' Addresses	23

1. Introduction

Network slicing allows a Service Provider to create independent and logical networks on top of a common or shared physical network infrastructure. Such network slices can be offered to customers or used internally by the Service Provider to facilitate or enhance their service offerings. A Service Provider can also use network slicing to structure and organize the elements of its infrastructure. [\[I-D.bestbar-teas-ns-packet\]](#) provides a path control technology agnostic solution that a Service Provider can deploy to realize network slicing in IP/MPLS networks. It introduces the notion of a slice aggregate and describes how a slice policy can be used to realize a slice aggregate by instantiating specific control and data plane behaviors on select topological elements in IP/MPLS networks.

In the control plane slice policy mode described in Section 4.2 of [\[I-D.bestbar-teas-ns-packet\]](#), the physical network resources in the network can be logically partitioned by having a representation of network resources appear in a virtual topology. The virtual topology can contain all or a subset of the physical network resource(s). The logical network resources that appear in the virtual topology can reflect a part, whole, or in- excess of the physical network resource capacity (when oversubscription is desirable). To perform network state dependent path computation and placement (Slice Aggregate aware TE) in this mode, the resource reservation on each link needs to be slice aggregate aware. Multiple slice policies may be applied on the same physical link. The slice aggregate network resource availability on links is updated when new paths are placed in the network. The slice aggregate resource reservation, in this case, can be maintained on each device or be centralized on a resource reservation manager that holds reservation states on links in the network.

This document will describe the IGP extensions required to describe a network slice aggregate's logical network resources.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. Terminology

The reader is expected to be familiar with the terminology specified in [[I-D.ietf-teas-ietf-network-slice-definition](#)], [[I-D.nsdt-teas-ns-framework](#)] and [[I-D.bestbar-teas-ns-packet](#)]. The term "Network Slice" used in this document must be interpreted as "IETF Network Slice" [[I-D.ietf-teas-ietf-network-slice-definition](#)]. For readability, a few key terms from these documents are repeated here:

IETF Network Slice: a well-defined composite of a set of endpoints, the Connectivity requirements between subsets of these endpoints, and Associated service requirements; the term 'network slice' in this Document refers to 'IETF network slice' [[I-D.ietf-teas-ietf-network-slice-definition](#)].

Slice Policy: a policy construct that enables instantiation of mechanisms in support of IETF network slice specific control and data plane behaviors on select topological elements; the enforcement of a slice policy results in the creation of a slice aggregate.

Slice Aggregate: a collection of packets that match a slice policy selection criteria and are given the same forwarding treatment; a slice aggregate comprises of one or more IETF network slice traffic streams; the mapping of one or more IETF network slices to a slice aggregate is maintained by the IETF Network Slice Controller.

Slice Aggregate aware TE (SA-TE): a mechanism for TE path selection that takes into account the available network resources associated with a specific slice aggregate.

4. Resource Allocation

For SA-TE, the existing "Maximum Reservable link bandwidth" MUST be maintained, but it remains generalized and interpreted as the aggregate usable bandwidth across all slice aggregates.

4.1. Resource Sharing and Resource Overbooking

[I-D.bestbar-teas-yang-ns-phd] allows a network administrator to apply different overbooking (or underbooking) ratios for different slice aggregates or groups of slice aggregates. The principal method to achieve this is through the definition of per slice aggregate maximum bandwidth value along with a Shared-Resource-Group identifier (SRG).

When the SRG is included for a slice aggregate within the slice policy, resources may be shared between slices of the same shared-resource-group. Strict resources sharing is only enforced between shared-resource-groups.

When there is resource contention, as a result of a reduction of network capacity, slice aggregate specific paths may be preempted on a priority basis to ensure continuous enforcement of the resource allocation policy as defined by the slice policy.

4.2. SA-TE Resource Allocation Model

[RFC4125] defines the Maximum Allocation Model (MAM) in support of [DST-TE REQ]. Resource guarantees, sharing and overbooking for slice aggregates follow the MAM.

5. Next-hop Filtering Capability

Per Slice Aggregate Next Hop filtering is an important capability for the path computing entity of a slice aggregate specific path. Consider the figure below and the incongruent slice aggregate membership outlined in Table 1.




Figure 1: Sample Topology

Link	IGP metric	Slice Aggregates
(r0,if1)-(r1,if1)	10	SA2, SA1
(r1,if2)-(r2,if1)	10	SA2
(r2,if2)-(r5,if1)	10	SA2
(r1,if3)-(r3,if1)	10	SA2
(r3,if3)-(r5,if2)	10	SA2
(r1,if4)-(r3,if2)	10	SA1
(r3,if4)-(r5,if3)	10	SA1
(r1,if5)-(r4,if1)	10	SA1
(r4,if2)-(r5,if4)	10	SA1
(r5,if5)-(r6,if1)	10	SA2, SA1

Table 1: Incongruent Slice Aggregate membership

Consider if r0 is the path computing entity and requires a segment-routing policy towards r6 via only links and nodes that are members of the SA2 slice aggregate and maximizing link utilization via ECMP.

r1 has 4 equal cost Next Hops to r6 over if2 (via r2), if3 (via r3), if4 (via r3), and if5 (via r4). However, the links from r1 over if4 (to r3) and over if5 (to r4) are not members of the SA2 slice.

Without any Per Slice Aggregate Next Hop filtering, r0 would need to compute two segment-lists consisting of several segments and load-balance equally between them to ensure traffic sent to r1 is forwarded over SA2 links. The two segment-lists may look like:

```
segment-list1: [<r2.Node-SID>, <r6.Node-SID>]
```

```
segment-list2: [<r1.if3.Adj-SID>, <r3.if3.Adj-SID>, <r6.Node-SID>]
```

When r1 supports Per Slice Aggregate Next Hop filtering on SR segments, it can advertise this capability in the SR network. r0 (or a PCE) can build a Per Slice aggregate topology and compute Per Slice aggregate ECMP paths for an SR node segment. Thus, for the same earlier example, r0 would be able to use a single segment (r6's node-SID) to steer over slice aggregate traffic over the SA2 slice.

As described in [[I-D.bestbar-spring-scalable-ns](#)], when a Global Identifier as Slice Selector is carried in the packet to identify the slice aggregate, r0 can use r6's node-SID to steer the packet over the specific next-hops that belong to SA2 slice aggregate.

For example on transit router r1, the SS field in the packet is used to identify the slice aggregate SA2, while r6.node-SID is used to

determine the subset of next-hops that are member of slice aggregate SA2.

6. IS-IS Extensions for SA-TE

As discussed in [[I-D.bestbar-teas-ns-packet](#)], routing protocols need to be extended to carry additional per slice link state and nodes need to advertise a slice policy capability.

This information will not be used by the IS-IS decision process.

6.1. Network Slicing GENINFO TLV

[RFC6823] describes the use of GENINFO TLV to advertise application-specific information that is not directly related to the operation of the IS-IS protocol. The current document defines a new Network Slicing Application Identifier under the Generic Information TLV (#251).

Both V and I flags of the Network Slicing GENINFO TLV SHOULD be unset.

The Network Slicing GENINFO TLV MAY be carried in the zero IS-IS instance or a non-zero IS-IS instance.

6.2. Slice Aggregate Aware Traffic Engineering Capabilities APPsub-TLV

This document defines a new APPsub-TLV under Network Slicing GENINFO TLV, to announce a node's SA-TE capability within a slice policy domain. A node can also advertise its ability to install slice aggregate specific Next Hops via this APPsub-TLV.

The Slice Aggregate Aware Traffic Engineering (SA-TE) Capabilities APPsub-TLV may contain optional sub-sub-TLVs. No sub-sub-TLVs are currently defined.

The SA-TE Capabilities APPsub-TLV has the following format:




Figure 2: SA-TE Capabilities APPsub-TLV

- o Type: (Value TBD by IANA)

- o Length: 2 + length of sub-sub-TLVs
- o Flags: 2 octets. The following flags are defined:

```

      0                               1
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+--+--+--+--+--+--+--+--+--+--+
|N|                                     |
+-+--+--+--+--+--+--+--+--+--+--+

```

Figure 3: Slicing Capability Flags

where:

N-flag: If set, the router supports Per Slice Next Hop filtering.

For purposes of extensibility, currently only one flag is defined. The receipt of a SA-TE Capabilities APPsub-TLV without N flag set is valid, but there is no semantic meaning defined at this time.

The SA-TE Capabilities APPsub-TLV MUST NOT appear more than once in a Network Slicing GENINFO TLV. If it appears more than once, the Network Slicing GENINFO TLV MUST be ignored by the receiver.

6.3. Slice Aggregate Aware Traffic Engineering Link APPsub-TLV

The SA-TE Link APPsub-TLV is an APPsub-TLV under the Network Slicing GENINFO TLV. This APPsub-TLV is used to advertise link information used by a SA-TE Application. It can carry link identification information which can be used by a SA-TE Application to uniquely identify a link.

The following illustrates encoding of the Value field of the IS-IS SA-TE Link APPsub-TLV.

	No. of octets
+-----+ System ID	6
+-----+ Psuedonode Number	1
+-----+ Flags	1
+-----+ Link Identification Information	0 to 24
+-----+ Sub-TLVs	Optional
+-----+	

Flags

```

      0 1 2 3 4 5 6 7
+---+---+---+---+
|Reserved |S|F|I|
+---+---+---+---+

```

Figure 4: SA-TE Link APPsub-TLV

The neighbor is identified by its System ID (6 octets), plus one octet to indicate the pseudonode number if the neighbor is on a LAN interface.

The following bit flags are defined

I bit (0x01): When the I bit is set, the 4-octet Link Local Identifier associated with the link immediately follows the Flags.

F bit (0x02): When the F bit is set, the 4-octet IPv4 interface address of the link immediately follows either the Flags (if I bit is clear) or the Link Local Identifier (if I bit is set).

S bit (0x04): When the S bit is set, the 16-octet IPv6 interface address of the link immediately follows the Flags (if both F and I bits are clear) or the Link Local Identifier (if I flag is set but F flag is not) or the IPv4 address (if F bit is set).

The Link Identification Information could have any or all of the Link Local Identifier, IPv4 interface address and IPv6 interface address of the link.

6.4. Slice Aggregate Aware Traffic Engineering Link Attributes Sub-sub-TLV

The SA-TE Link Attributes Sub-sub-TLV is a sub-sub-TLV under the SA-TE Link APPsub-TLV. This sub-sub-TLV advertises various bandwidth attributes on a particular link for a specific network slice.

There MAY be a SA-TE Link Attributes Sub-sub-TLV for each network slice the link participates in.

The "SA-TE Link Attributes" sub-sub-TLV format is illustrated below:

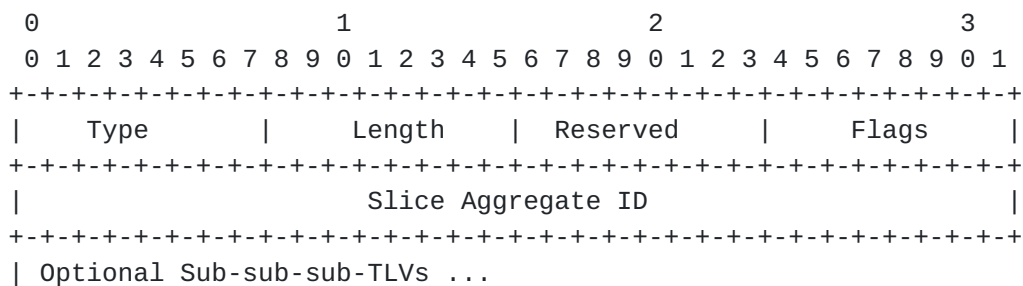


Figure 5: SA-TE Link Attributes Sub-sub-TLV

Type: (Value TBD by IANA)

Length: 4 to 20 based on Flags

Slice Aggregate ID: 32-bit slice aggregate identifier.

Flags: 1-octet field for future use.

6.4.1. Slice Aggregate Aware Traffic Engineering Unreserved Bandwidth Sub-sub-sub-TLV

The SA-TE Unreserved Bandwidth Sub-sub-sub-TLV is a Sub-sub-sub-TLV under the SA-TE Link Attributes Sub-sub-TLV. This Sub-sub-sub-TLV advertises the unreserved bandwidth on a particular link for a specific network slice aggregate.

The SA-TE Unreserved Bandwidth Sub-sub-sub-TLV can carry up to eight bandwidth values, one for each priority level. The units are bytes (not bits!) per second.

The "SA-TE Unreserved Bandwidth" Sub-sub-sub-TLV format is illustrated below:

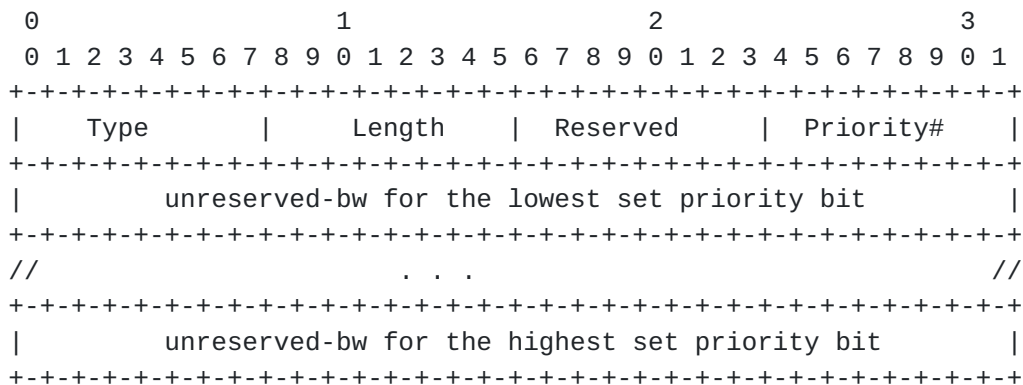


Figure 6: SA-TE Unreserved Bandwidth Sub-sub-sub-TLV

- o Type: (Value TBD by IANA)
- o Length: Variable.
- o Priority#: Bit-map defining the priority values in use. Up to 8 priority values can be used.
- o Unreserved-bw: Unreserved link bandwidth on the link where each unreserved-bw entry corresponds to a non-zero bit in the Priority field, starting at the least significant bit.

6.4.2. Slice Aggregate Aware Traffic Engineering Residual Bandwidth Sub-sub-sub-TLV

The SA-TE Residual Bandwidth Sub-sub-sub-TLV is a Sub-sub-sub-TLV under the SA-TE Link Attributes Sub-sub-TLV. This Sub-sub-sub-TLV advertises the residual bandwidth on a particular link for a specific Network slice aggregate.

The format of SA-TE Residual Bandwidth Sub-sub-sub-TLV matches the format of Unidirectional Residual Bandwidth Sub-TLV defined in [\[RFC8570\]](#).

6.4.3. Slice Aggregate Aware Traffic Engineering Available Bandwidth Sub-sub-sub-TLV

The SA-TE Available Bandwidth Sub-sub-sub-TLV is a Sub-sub-sub-TLV under the SA-TE Link Attributes Sub-sub-TLV. This Sub-sub-sub-TLV advertises the available bandwidth on a particular link for a specific Network slice aggregate.

The format of SA-TE Available Bandwidth Sub-sub-sub-TLV matches the format of Unidirectional Available Bandwidth Sub-TLV defined in [\[RFC8570\]](#).

6.4.4. Slice Aggregate Aware Traffic Engineering Utilized Bandwidth Sub-sub-sub-TLV

The SA-TE Utilized Bandwidth Sub-sub-sub-TLV is a Sub-sub-sub-TLV under the SA-TE Link Attributes Sub-sub-TLV. This Sub-sub-sub-TLV advertises the utilized bandwidth on a particular link for a specific Network slice aggregate.

The format of SA-TE Utilized Bandwidth Sub-sub-sub-TLV matches the format of Unidirectional Utilized Bandwidth Sub-TLV defined in [\[RFC8570\]](#).

7. OSPF Extensions for SA-TE

A future version of this document will define OSPF extensions for SA-TE.

8. Functional Example of Unreserved-Bandwidth Advertisements

For illustration purposes, we now present a few examples of how these extensions may be used.

8.1. Formula to Compute the Unreserved Bandwidth

For a slice 's' at priority 'p' on a specific TE link, where:

s : Slice-name

p : reservation priority (0-7)

{ S } : Set containing all slice names in the network

$$\text{Unreserv-BW}(s, p) = \text{MIN} ($$

$$[\text{Max-resv-shared-BW}(s) - \text{Sum} \{ \text{Reserved-BW}(s, q) \}] \text{ for all } q \leq p,$$

$$[\text{Max-resv-BW}(\text{interface}) - \text{Sum} \{ \text{Reserved-BW}(c, q) \}] \text{ for all } c =$$

$$\{ S \} \ \& \ q \leq p$$

$$)$$

The formula above is generic and applies for all priority values (0-7)

8.2. Examples

All these examples assume two slice aggregates, SA2 and SA1, with the following configuration parameters:

Per TE link resource allocation:

Max-BW(interface): 10G (derived from the physical interface BW)

Max-resv-BW(interface): 10G (default: equal to Max-BW)

Slice Policy:

Max-resv-shared-BW(SA1): 10G

Max-resv-shared-BW(SA2): 7G

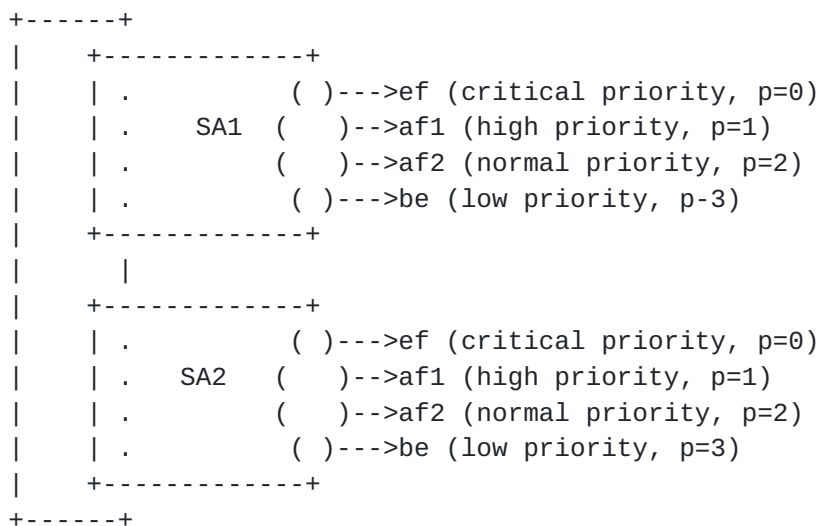


Figure 7: Example 10G Interface with 2 Slice Aggregates

The reservation priority in the following examples will use the names critical, high, normal and low corresponding to priority values 0-3. The unreserved-BW for each slice, priority can be described with the following formulas:

Unreserv-BW(SA1, critical) = MIN (

[Max-resv-shared-BW(SA1) - Sum { Reserved-BW(SA1, q) }] for q = { critical },

[Max-resv-BW(interface) - Sum { Reserved-BW(c, q) }] for c = { SA1, SA2 } & q = { critical }

)

Unreserv-BW(SA1, high) = MIN (

[Max-resv-shared-BW(SA1) - Sum { Reserved-BW(SA1, q) }] for q =
{ critical, high },

[Max-resv-BW(interface) - Sum { Reserved-BW(c, q) }] for c = {
SA1, SA2 } & q = { critical, high }

)

Unreserv-BW(SA1, normal) = MIN (

[Max-resv-shared-BW(SA1) - Sum { Reserved-BW(SA1, q) }] for q =
{ critical, high, normal },

[Max-resv-BW(interface) - Sum { Reserved-BW(c, q) }] for c = {
SA1, SA2 } & q = { critical, high, normal }

)

Unreserv-BW(SA1, low) = MIN (

[Max-resv-shared-BW(SA1) - Sum { Reserved-BW(SA1, q) }] for q =
{ critical, high, normal, low },

[Max-resv-BW(interface) - Sum { Reserved-BW(c, q) }] for c = {
SA1, SA2 } & q = { critical, high, normal, low }

)

Unreserv-BW(SA2, critical) = MIN (

[Max-resv-shared-BW(SA2) - Sum { Reserved-BW(SA2, q) }] for q =
{ critical },

[Max-resv-BW(interface) - Sum { Reserved-BW(c, q) }] for c = {
SA1, SA2 } & q = { critical }

)

Unreserv-BW(SA2, high) = MIN (

[Max-resv-shared-BW(SA2) - Sum { Reserved-BW(SA2, q) }] for q =
{ critical, high },


```

    [ Max-resv-BW(interface) - Sum { Reserved-BW(c, q) } ] for c = {
    SA1, SA2 } & q = { critical, high }

)

Unreserv-BW(SA2, normal) = MIN (

    [ Max-resv-shared-BW(SA2) - Sum { Reserved-BW(SA2, q) } ] for q =
    { critical, high, normal },

    [ Max-resv-BW(interface) - Sum { Reserved-BW(c, q) } ] for c = {
    SA1, SA2 } & q = { critical, high, normal }

)

Unreserv-BW(SA2, low) = MIN (

    [ Max-resv-shared-BW(SA2) - Sum { Reserved-BW(SA2, q) } ] for q =
    { critical, high, normal, low },

    [ Max-resv-BW(interface) - Sum { Reserved-BW(c, q) } ] for c = {
    SA1, SA2 } & q = { critical, high, normal, low }

)

```

8.2.1. Example 1 Results

Time0: no LSPs

Unreserv-BW(SA1, q) = 10G, for q = { critical, high, normal, low }

Unreserv-BW(SA2, q) = 10G, for q = { critical, high, normal, low }

Time1: new SA1, normal LSP 5G

Unreserv-BW(SA1, normal) = 5G

Unreserv-BW(SA1, low) = 5G

Unreserv-BW(SA2, normal) = 5G

Unreserv-BW(SA1, low) = 5G

No change to other Unreserv-BW values

Note: If a reservation is made at a priority p, then there will be no change to unreserv-bw(s, q) for all q < p.

Time2: new SA1, critical LSP 5G

Unreserv-BW(SA1, critical) = 5G

Unreserv-BW(SA1, high) = 5G

Unreserv-BW(SA1, normal) = 0G

Unreserv-BW(SA1, low) = 0G

Unreserv-BW(SA2, critical) = 5G

Unreserv-BW(SA2, high) = 5G

Unreserv-BW(SA2, normal) = 0G

Unreserv-BW(SA2, low) = 0G

Time3: new SA2, critical LSP 5G

-> preempt LSP(SA1, normal)

Unreserv-BW(SA1, critical) = 0G

Unreserv-BW(SA1, high) = 0G

Unreserv-BW(SA2, critical) = 0G

Unreserv-BW(SA2, high) = 0G

No change to other Unreserv-BW values

8.2.2. Example 2 Results

Time0: no LSPs

Unreserv-BW(SA1, q) = 10G, for q = { critical, high, normal, low }

Unreserv-BW(SA2, q) = 10G, for q = { critical, high, normal, low }

Time1: new SA1, critical LSP 5G

Unreserv-BW(SA1, critical) = 5G

Unreserv-BW(SA1, high) = 5G

Unreserv-BW(SA1, normal) = 5G

Unreserv-BW(SA1, low) = 5G

Unreserv-BW(SA2, critical) = 5G

Unreserv-BW(SA2, high) = 5G

Unreserv-BW(SA2, normal) = 5G

Unreserv-BW(SA2, low) = 5G

Time2: new SA1, critical LSP 5G

Unreserv-BW(SA1, critical) = 0G

Unreserv-BW(SA1, high) = 0G

Unreserv-BW(SA1, normal) = 0G

Unreserv-BW(SA1, low) = 0G

Unreserv-BW(SA2, critical) = 0G

Unreserv-BW(SA2, high) = 0G

Unreserv-BW(SA2, normal) = 0G

Unreserv-BW(SA2, low) = 0G

8.2.3. Example 3 Results

Time0: no LSPs

Unreserv-BW(SA1, q) = 10G, for q = { critical, high, normal, low }

Unreserv-BW(SA2, q) = 10G, for q = { critical, high, normal, low }

Time1: new SA2, critical LSP 10G

Unreserv-BW(SA1, critical) = 0G

Unreserv-BW(SA1, high) = 0G

Unreserv-BW(SA1, normal) = 0G

Unreserv-BW(SA1, low) = 0G

Unreserv-BW(SA2, critical) = 0G

Unreserv-BW(SA2, high) = 0G

Unreserv-BW(SA2, normal) = 0G

Unreserv-BW(SA2, low) = 0G

Time2: new SA1, normal LSP 5G

-> not admitted

9. Scaling Considerations

9.1. Link-State Update Frequency

SA-TE applications SHOULD rate limit the link attribute updates notified to IGP.

10. IANA Considerations

10.1. Network Slicing GENINFO TLV

IANA is requested to allocate the IS-IS Application Identifier TBD [2 suggested] under the Generic Information TLV (#251) [[RFC6823](#)] for Network Slicing.

This document defines APPSub-TLVs under Network Slicing GENINFO TLV, for which IANA is requested to create a new registry entitled "Network Slicing GENINFO Parameters".

10.2. Network Slicing GENINFO APPsub-TLVs

IANA is requested to create a subregistry in the Network Slicing GENINFO Parameters Registry as follows:

Sub-Registry: Network Slicing APPsub-TLV Types under IS-IS TLV #251 Application Identifier #TBD

Registration Procedures: IETF Review with additional requirements on the documentation of the use being registered as specified in [Section 10.2](#) of this document.

Reference: this document

Type	Name	Reference
-----	-----	-----
0	Reserved	<this document>
1	SA-TE Capabilities	<this document>
2	SA-TE Link	<this document>
3-254	Unassigned	<this document>

Network Slicing APPsub-TLV Types 3 through 254 are available for assignment by IETF Review. The RFC causing such an assignment will also include a discussion of security issues and of the rate of change of the information being advertised. Network Slicing APPsub-TLVs MUST NOT alter basic IS-IS protocol operation including the establishment of adjacencies and the decision process for IS-IS [IS-IS], [[RFC1195](#)].

10.3. SA-TE Link Sub-sub-TLVs

IANA is requested to create a subregistry in the Network Slicing GENINFO Parameters Registry as follows:

Sub-Registry: SA-TE Sub-sub-TLV Types under SA-TE Link APPsub-TLV #TBD

Registration Procedures: IETF Review with additional requirements on the documentation of the use being registered as specified in [Section 10.3](#) of this document.

Reference: this document

Type	Name	Reference
-----	-----	-----
0	Reserved	<this document>
1	SA-TE Link Attributes	<this document>
2-254	Unassigned	<this document>

Types 2 through 254 are available for assignment by IETF Review. The RFC causing such an assignment will also include a discussion of security issues and of the rate of change of the information being advertised.

10.4. SA-TE Link Attributes Sub-sub-sub-TLVs

IANA is requested to create a subregistry in the Network Slicing GENINFO Parameters Registry as follows:

Sub-Registry: SA-TE Sub-sub-sub-TLV Types under SA-TE Link Attributes Sub-sub-TLV #TBD

Registration Procedures: IETF Review with additional requirements on the documentation of the use being registered as specified in [Section 10.4](#) of this document.

Reference: this document

Type	Name	Reference
-----	-----	-----
0	Reserved	<this document>
1	SA-TE Unreserved Bandwidth	<this document>
2-36	Unassigned	<this document>
37	SA-TE Residual Bandwidth	<this document>
38	SA-TE Available Bandwidth	<this document>
39	SA-TE Utilized Bandwidth	<this document>
2-254	Unassigned	<this document>

Types 2 through 36 and 40 through 254 are available for assignment by IETF Review. The RFC causing such an assignment will also include a discussion of security issues and of the rate of change of the information being advertised. Type 37, 38 and 39 have been suggested to match with the Sub-TLVs 37, 38 and 39 mentioned in the IS-IS "Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223" Registry.

[11.](#) Security Considerations

The advertisement of the Network Slicing GENINFO TLV and its APPsub-TLVs raises no new security issues for IS-IS. This information will not be used by the IS-IS decision process. Where appropriate, it is recommended that either [[RFC5304](#)] or [[RFC5310](#)] be used.

[12.](#) References

[12.1.](#) Normative References

[I-D.bestbar-spring-scalable-ns]

Saad, T. and V. Beeram, "Scalable Network Slicing over SR Networks", [draft-bestbar-spring-scalable-ns-00](#) (work in progress), December 2020.

[I-D.bestbar-teas-ns-packet]

Saad, T., Beeram, V., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., and X. Liu, "Realizing Network Slices in IP/MPLS Networks", [draft-bestbar-teas-ns-packet-01](#) (work in progress), December 2020.

[I-D.bestbar-teas-yang-ns-phd]

Saad, T. and V. Beeram, "YANG Data Model for Network Slice Per-Hop Definition", [draft-bestbar-teas-yang-ns-phd-00](#) (work in progress), November 2020.

[I-D.ietf-teas-ietf-network-slice-definition]

Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "Definition of IETF Network Slices", [draft-ietf-teas-ietf-network-slice-definition-00](#) (work in progress), January 2021.

[I-D.nsdt-teas-ns-framework]

Gray, E. and J. Drake, "Framework for Transport Network Slices", [draft-nsdt-teas-ns-framework-04](#) (work in progress), July 2020.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.

[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>>.

[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>>.

[RFC4124] Le Faucheur, F., Ed., "Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering", [RFC 4124](#), DOI 10.17487/RFC4124, June 2005, <<https://www.rfc-editor.org/info/rfc4124>>.

[RFC4125] Le Faucheur, F. and W. Lai, "Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering", [RFC 4125](#), DOI 10.17487/RFC4125, June 2005, <<https://www.rfc-editor.org/info/rfc4125>>.

[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>>.

[RFC6823] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", [RFC 6823](#), DOI 10.17487/RFC6823, December 2012, <<https://www.rfc-editor.org/info/rfc6823>>.

- [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>>.
- [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>>.

[12.2](#). Informative References

- [ISO10589] IANA, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", August 1987, <ISO/IEC 10589:2002>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", [RFC 7981](#), DOI 10.17487/RFC7981, October 2016, <<https://www.rfc-editor.org/info/rfc7981>>.
- [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>>.
- [RFC8202] Ginsberg, L., Previdi, S., and W. Henderickx, "IS-IS Multi-Instance", [RFC 8202](#), DOI 10.17487/RFC8202, June 2017, <<https://www.rfc-editor.org/info/rfc8202>>.
- [RFC8919] Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS Application-Specific Link Attributes", [RFC 8919](#), DOI 10.17487/RFC8919, October 2020, <<https://www.rfc-editor.org/info/rfc8919>>.

[Appendix A](#). Contributors

Srihari Sangli
Juniper Networks
Email: ssangli@juniper.net

Vishnu Pavan Beeram
Juniper Networks
Email: vbeeram@juniper.net

Tarek Saad
Juniper Networks
Email: tsaad@juniper.net

Chandra Ramachandran
Juniper Networks
Email: csekar@juniper.net

[Appendix B](#). Acknowledgements

The authors would like to thank Chris Bowers and Shraddha Hegde for their valuable review and feedback.

Authors' Addresses

William Britto
Juniper Networks

Email: bwilliam@juniper.net

Rajesh Shetty
Juniper Networks

Email: mrajesh@juniper.net

Colby Barth
Juniper Networks

Email: cbarth@juniper.net

Bin Wen
Comcast

Email: Bin_Wen@cable.comcast.com

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn