

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

T. Saad
V. Beeram
Juniper Networks
R. Chen
S. Peng
ZTE Corporation
B. Wen
Comcast
D. Ceccarelli
Ericsson
February 22, 2021

Scalable Network Slicing over SR Networks
draft-bestbar-spring-scalable-ns-01

Abstract

Multiple network slices can be realized on top of a single shared network. A router that requires forwarding of a packet that belongs to a slice aggregate may have to decide on the forwarding action to take based on selected next-hop(s), and the forwarding treatment (e.g., scheduling and drop policy) to enforce based on the slice aggregate per-hop behavior. Segment Routing is a technology that enables the steering of packets in a network by encoding pre-established segments within the network into the packet header. This document introduces mechanisms to enable forwarding of packets over a specific slice aggregate along a Segment Routing (SR) path.

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.

Internet-Draft

Scalable Network Slices over SR

February 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	2
1.1.	Terminology	3
2.	Forwarding over SR Network Slices	3
2.1.	Path Selection	4
2.2.	Network Slice Selection	4
2.2.1.	Segment Range as Slice Selector	4
2.2.2.	Global Identifier as Slice Selector	7
3.	IANA Considerations	9
4.	Security Considerations	9
5.	Acknowledgement	10
6.	Contributors	10
7.	References	10
7.1.	Normative References	10
7.2.	Informative References	12
	Authors' Addresses	12

[1.](#) Introduction

Network slicing allows a service provider, or a network operator to create independent and isolated logical networks on top of a common or shared physical network infrastructure.

When logical network slices are realized on top of a shared physical network, it is important to forward traffic using only the specific resource(s) allocated to the network slice.

The definition of a network slice for use within the IETF and the characteristics of IETF network slice are specified in [\[I-D.nsdt-teas-ietf-network-slice-definition\]](#). A framework for reusing IETF VPN and traffic-engineering technologies to realize IETF network slices is discussed in [\[I-D.nsdt-teas-ns-framework\]](#).

[\[I-D.bestbar-teas-ns-packet\]](#) introduces the notion of a Slice Aggregate as the construct that comprises of one or more IETF network slice traffic streams. A slice policy can be used to realize a slice aggregate by instantiating specific control and data plane resources on select topological elements in an IP/MPLS network. The packets belonging to a specific slice aggregate MAY require to be identified so that a specific forwarding treatment (e.g., scheduling and drop policy) is enforced.

Segment Routing (SR) [\[RFC8402\]](#) specifies a mechanism to steer packets through a network by carrying an ordered list of segments. A segment is referred to by its Segment Identifier (SID).

This document introduces two approaches applicable to SR networks that enable forwarding of packet(s) that belong to a slice aggregate over a SR Path.

The first approach extends the SR paradigm by defining a new SID type (slice SID) that, in addition to defining the forwarding action (next-hop selection), associates a SID to slice aggregate and allows enforcing the associated forwarding treatment. The extensions to IGP for slice aggregate SIDs are defined in [\[I-D.bestbar-lsr-spring-sa\]](#).

The second approach relies on a separate field that is carried in the packet (e.g., MPLS label) to identify the slice aggregate and uses another field (e.g., existing SR segments) for the path selection for the packet.

[1.1](#). Terminology

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

[2.](#) Forwarding over SR Network Slices

A router that receives a packet that belongs to a slice aggregate has to decide on the set of eligible next-hop(s) to forward the packet on (path selection), and on the forwarding treatment (scheduling and drop policy) that needs to be enforced for a specific slice aggregate (slice aggregate selection).

[2.1.](#) Path Selection

The segment routing architecture [[RFC8402](#)] defines a number of topological segments that may be advertised in routing protocols to allow for a flexible definition of end-to-end paths. For example, an SR-capable IGP router may advertise SIDs for its attached prefixes and adjacencies.

The IGP-Adjacency segment represents the strict path over a specific adjacency between two routers, while the IGP-Prefix segment represents the path to a prefix that is computed over a specific topology and algorithm.

For an IGP-Prefix segment, the IGP uses the topology and algorithm to compute the primary, and optionally alternate (backup) next-hop(s) for a destination prefix. SR allows the use of multiple routing algorithms (e.g., Flexible Algorithms) that enable IGPs on a router to compute paths for Prefix-SIDs whose topology may be constrained and whose paths optimized for additional metric types other than the default IGP cost (e.g., delay metric).

Multiple slice aggregates may overlap over the same topology and require paths for prefixes to be optimized for the same Algorithm. In such case, the IGP selected path for the slice aggregate Prefix-SIDs can share the same IGP computed path (including the primary and backup next-hop(s)). This enables the IGP to optimize the path computation and path programming for such SA Prefix-SIDs.

[2.2.](#) Network Slice Selection

The routers in network that forward traffic over links that are shared by multiple slice aggregates need to identify the slice aggregate that the packet belongs to in order to enforce the associated forwarding treatment on it.

[I-D.bestbar-teas-ns-packet] introduces the slice policy as a means to realize a slice aggregate by instantiating specific control and/or data plane resources on select topological elements in the network. In order to enforce a forwarding treatment associated with a slice aggregate, the packets traversing a router MUST be identified as part of a slice aggregate (for example, by using field(s) carried in the packet).

[2.2.1.](#) Segment Range as Slice Selector

It is possible to derive the forwarding action (next-hop selection) and the per-hop forwarding treatment from a single field (e.g. SR segment) carried inside a packet that is traversing the SR network.

For example, one way to achieve this is leverage the SR Flexible-Algorithm [[I-D.ietf-lsr-flex-algo](#)] to assign SR SID per slice aggregate. A router can assign and advertise SR Prefix-SIDs per Flex-Algorithm for a prefix to enable reachability over multiple slice aggregates.

For a specific Flexible Algorithm, the range of Prefix-SIDs of all prefixes in the network can be used as a slice selector mapping to a specific slice aggregate. This approach does not require protocol extensions to be realized; however, it poses serious IGP scalability concerns when realizing a large number of slice aggregates.

Alternatively, this document proposes to extend the IGP SR Prefix-SID and Adjacency-SID sub-TLVs defined in [[RFC8667](#)] and [[RFC8665](#)] to carry an additional distinguisher (Slice Aggregate identifier) to allow multiple SIDs to be assigned (and advertised) for the same topological element for the same Flexible-Algorithm and topology. In such a case, a transit router can use the SR slice aggregate SID carried in the packet to identify the slice aggregate, as well as to determine the forwarding next-hop.

Multiple Slice Aggregate Prefix-SIDs (SA Prefix-SIDs) can be assigned

to the same prefix, while they share the same topology and Algorithm. The SA Prefix-SIDs can also share the same IGP computed paths (primary and backup). Similarly, multiple Slice Aggregate Adjacency-SIDs (SA Adjacency-SIDs) can be allocated for the same adjacency between the two routers to distinguish forwarding over the same adjacency for each slice aggregate. The extensions for IGPs to advertise SA Prefix-SIDs and SA Adjacency-SIDs are defined in [\[I-D.bestbar-lsr-spring-sa\]](#).

The same forwarding treatment MUST be enforced on all packets belonging to a slice aggregate but destined to different topological elements in the network. In this case, a range of SA (Prefix and Adjacency) SIDs is used to select the slice aggregate, and hence enforce the same forwarding treatment on them.

Note that this approach requires maintaining per slice aggregate state for each topological element on every router in the network in both the control and data plane. For example, a network composed of 'N' routers, where each router has up to 'K' adjacencies to its neighbors, a router would have to assign and advertise 'M' SA Prefix-SIDs and 'M' SA Slice Adjacency-SID(s) to each of its 'K' adjacencies. Consequently, a router will have to maintain up to $(N+K)*M$ SIDs in the control plane, and an equal number of labeled routes in its forwarding plane.

Consider a network shown in Figure 1 that is enabled for SR. The Segment Routing Global Block (SRGB) on all routers is assumed to start from 16000. We assume the links in the network are partitioned amongst two network slice aggregates: SA1, and SA2.

- o Node R5 assigns two Algorithm 0 SA Prefix-SIDs, index=105, and index=205 to represent the shortest IGP to R5 for slice aggregates SA1 and SA2 respectively.
- o A Flexible Algorithm Definition (FAD) for Algorithm 128 is defined by the user such that the FAD Metric-Type is 1 (Min Unidirectional Link Delay).
- o Node R5 assigns two Algorithm 128 SA Prefix-SIDs, index=815, and index=825 to represent the least delay path to R5 for slice

aggregates SA1 and SA2 respectively.

- o All routers in the network participate and advertise their capability to compute FAD 128 Prefix-SID paths.

Using the approach described in this section, R1 is able to forward packets that traverse slices aggregates SA1 and SA2 along the least delay path by imposing the MPLS SR SID 16815, and 16825 respectively.

In addition, R1 is able to forward packets that traverse slice aggregate SA1 and SA2 along the IGP shortest path to R5 by imposing the MPLS SR SID 16015, and 16025 respectively.

SLICE	ALG(0)	Path
AGG	SA Prefix-SID(R5)	Symbol
=====		
SA1	16015	+
SA2	16025	@
..		
SAn	..	

SLICE	ALG(128)	Path
AGG	SA Prefix-SID(R5)	Symbol
=====		
SA1	16815	.
SA2	16825	*
..		
SAn	..	




Figure 1: Example of forwarding over slice aggregates using SR Paths.

[2.2.2.](#) Global Identifier as Slice Selector

It is possible that the forwarding action and the per-hop behavior treatment is derived from different fields carried in a packet. For example, a packet can carry a global slice selector field that can be used to define the forwarding treatment while the forwarding next-hop relies on the SR topological SIDs. This makes the slice aggregate identification independent of the topology or the destination of the packet, and thus, allows for scalable slice aggregates.

The Slice aggregate Selector (SS) is carried in each packet destined to any topological element and that is to be steered over the slice aggregate. For example, the slice aggregate SS can be carried in an MPLS label that is present in an MPLS packet's label stack. It is

possible, also, to have a range of MPLS labels to represent the SS

associated with slice aggregate.

When the slice aggregate is realized over an IPv6 dataplane, the SS can be encoded in the IP header. For example, the SS can be encoded in a portion of the IPv6 Flow Label field as described in [\[I-D.filsfils-spring-srv6-stateless-slice-id\]](#).

Routers within the network use the topological SR segment SIDs to determine the forwarding action (next-hop selection), and use the slice aggregate selector to enforce the dataplane policy (e.g., as defined by the slice policy in [\[I-D.bestbar-teas-ns-packet\]](#)).

The SS label may be embedded at different positions in the MPLS label stack. For example, the SS label MAY be located at the top of the MPLS packet label stack and maintained, by each hop, while the packet is forwarded along the SR path. However, since assigning a global MPLS label on all nodes for the SS may not be always feasible, an alternative is to assign a global Index for a Slice Aggregate Selector (SA Selector Index). In this case, the SA Selector Index is used to determine the actual MPLS label value (e.g., from the router Global Label Block) on a given router.

The SS label can also reside at the bottom of the label stack. For example, a range of VPN service labels may also serve as a SS to map traffic from multiple VPNs to the same slice aggregate.

Another option is to encode the SS as part of a well-known label such as Entropy Label (EL) as suggested in [\[I-D.decreaene-mpls-slid-encoded-entropy-label-id\]](#). This optimizes the number of the MPLS labels needed in the stack and provides an ease incremental deployment.

Lastly, a new Special Purpose Label- e.g., Slice Selector Indicator (SSI)- from the MPLS the Base Special-Purpose MPLS Label, or Extended Special-Purpose MPLS Label spaces can be used to indicate that a SS label immediately follows the SSI. In this case, the ingress router of slice aggregate boundary will impose at least two additional MPLS labels (SSI + SS) to identify the packets that belong to the slice aggregate.

This approach reduces the amount of state required to be stored on a router to allow forwarding over slice aggregates since it does not require a Prefix-SID state per slice aggregate in the control plane, nor in the forwarding plane.

To illustrate forwarding over slice aggregates using a SS label, we consider the same network described earlier in Figure 1, but with some changes in the configuration:

- o Node R5 assigns an Algorithm 0 Prefix-SID of index=5 to represent the shortest IGP path from any router to R5.
- o Node R5 assigns Algorithm 128 Prefix-SID of index=805 to represent the least delay path from any router to R5.
- o All routers in the network participate and advertise their capability to compute FAD 128 Prefix-SID paths.
- o The SS labels 1001 and 1002 are used for packets that require to be forwarded over slice aggregates SA1 and SA2 respectively.

Using the approach described in this section, R1 is able to forward packets that traverse slice aggregate SA1 and SA2 along the least delay path by imposing the following labels {16805, SSI, 1001} and {16805, SSI, 1002} respectively.

Similarly, R1 is able to forward packets that traverse over slice aggregates SA1 and SA2 along the IGP shortest path to R5 by imposing the following labels {16005, SSI, 1001} and {16005, SSI, 1002} respectively. The path that the packets traverse in each of the above case remains as described in Figure 1.

[3.](#) IANA Considerations

This document has no IANA actions.

[4.](#) Security Considerations

The main goal of network slicing is to allow for some level of isolation for traffic from multiple different network slices that are utilizing a common network infrastructure and to allow for different levels of services to be provided for traffic traversing a single slice aggregate resource(s).

A variety of techniques may be used to achieve this, but the end result will be that some packets may be mapped to specific resource(s) and may receive different (e.g., better) service treatment than others. The mapping of network traffic to a specific slice is indicated primarily by the SS, and hence an adversary may be able to utilize resource(s) allocated to a specific network slice by injecting packets carrying the same SS field in their packets.

Such theft-of-service may become a denial-of-service attack when the modified or injected traffic depletes the resources available to forward legitimate traffic belonging to a specific slice aggregate.

The defense against this type of theft and denial-of-service attacks consists of the combination of traffic conditioning at network slicing domain boundaries with security and integrity of the network infrastructure within a network slicing domain.

[5.](#) Acknowledgement

The authors would like to thank Krzysztof Szarkowicz, Swamy SRK, Navaneetha Krishnan, and Prabhu Raj Villadathu Karunakaran for their review of this document, and for providing valuable feedback on it.

[6.](#) Contributors

The following individuals contributed to this document:

Colby Barth
Juniper Networks
Email: cbarth@juniper.net

Srihari R. Sangli
Juniper Networks
Email: ssangli@juniper.net

Chandra Ramachandran
Juniper Networks
Email: csekar@juniper.net

[7.](#) References

[7.1.](#) Normative References

- [I-D.bestbar-lsr-spring-sa]
Saad, T., Beeram, V., Chen, R., Peng, S., Wen, B., and D. Ceccarelli, "IGP Extensions for SR Slice Aggregate SIDs", February 2021.

[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.

Saad, et al.

Expires August 26, 2021

[Page 10]

Internet-Draft

Scalable Network Slices over SR

February 2021

[I-D.decraene-mpls-slid-encoded-entropy-label-id]

Decraene, B., Filsfils, C., Henderickx, W., Saad, T., and V. Beeram, "Using Entropy Label for Network Slice Identification in MPLS networks.", [draft-decraene-mpls-slid-encoded-entropy-label-id-00](#) (work in progress), December 2020.

[I-D.filsfils-spring-srv6-stateless-slice-id]

Filsfils, C., Clad, F., Camarillo, P., and K. Raza, "Stateless and Scalable Network Slice Identification for SRv6", [draft-filsfils-spring-srv6-stateless-slice-id-02](#) (work in progress), January 2021.

[I-D.ietf-lsr-flex-algo]

Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-13](#) (work in progress), October 2020.

[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>>.

- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [RFC 8665](#), DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [RFC 8667](#), DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

Saad, et al.

Expires August 26, 2021

[Page 11]

Internet-Draft

Scalable Network Slices over SR

February 2021

[7.2.](#) Informative References

- [I-D.nsd-t-eas-ietf-network-slice-definition]
Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "Definition of IETF Network Slices", [draft-nsdt-teas-ietf-network-slice-definition-02](#) (work in progress), December 2020.
- [I-D.nsd-t-eas-ns-framework]
Gray, E. and J. Drake, "Framework for Transport Network Slices", [draft-nsdt-teas-ns-framework-04](#) (work in progress), July 2020.

Authors' Addresses

Tarek Saad
Juniper Networks

Email: tsaad@juniper.net

Vishnu Pavan Beeram
Juniper Networks

Email: vbeeram@juniper.net

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Bin Wen
Comcast

Email: Bin_Wen@cable.comcast.com

Saad, et al.

Expires August 26, 2021

[Page 12]

Internet-Draft

Scalable Network Slices over SR

February 2021

Daniele Ceccarelli
Ericsson

Email: daniele.ceccarelli@ericsson.com

