

Workgroup: SPRING
Internet-Draft:
draft-filsfils-spring-srv6-stateless-slice-
id-04
Published: 30 July 2021
Intended Status: Standards Track
Expires: 31 January 2022
Authors: C. Filsfils, Ed. F. Clad, Ed.
 Cisco Systems, Inc. Cisco Systems, Inc.
 P. Camarillo K. Raza
 Cisco Systems, Inc. Cisco Systems, Inc.
 D. Voyer R. Rokui
 Bell Canada Nokia

Stateless and Scalable Network Slice Identification for SRv6

Abstract

This document defines a stateless and scalable solution to achieve network slicing with SRv6.

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 31 January 2022.

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. Slice Identifier](#)
- [3. SLID Presence Indicator](#)
- [4. Ingress PE SLID Assignment](#)
- [5. Per-Slice Forwarding](#)
- [6. Bandwidth-Allocation Slice](#)
- [7. Backward Compatibility](#)
- [8. Acknowledgements](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

SRv6 Network Programming[[RFC8986](#)] enables the creation of overlays with underlay optimization to be deployed in an SR domain[[RFC8402](#)].

As defined in [[RFC8754](#)], all inter-domain packets are encapsulated for the part of the packet journey that is within the SR domain. The outer IPv6 header is originated by a node of the SR domain and is destined to a node of the SR domain.

This document describes a stateless encoding of slice identification in the outer IPv6 header of an SR domain. The slice identification is independent of topology and the QoS/DiffServ policy of the network, thus enabling scalable network slicing for SRv6 overlays.

The definition of network slicing in the context of networks built from IETF technologies is specified in [[I-D.ietf-teas-ietf-network-slices](#)]. It defines the term "IETF Network Slice" and establishes the general principles of network slicing in the IETF context. It also discusses the general framework for requesting and operating IETF Network Slices, the characteristics of an IETF Network Slice, the necessary system components and interfaces, and how abstract requests can be mapped to more specific technologies. The document also discusses related considerations with monitoring and security.

2. Slice Identifier

The Slice Identifier (SLID) is a value encoded within the IPv6 packet that allows transit routers to process the packet according to network slice-based policy. An example of slice-based policy that can be enforced using the SLID is described in [Section 6](#).

The SLID may identify a unique IETF network slice or a group of slices that share the same policy. For example, a SLID may identify a slice aggregate [[I-D.bestbar-teas-ns-packet](#)].

This document proposes to encode the SLID in a portion of the IPv6 Flow Label.

The precise SLID location within the IPv6 Flow Label and the number of bits used to encode it are governed by local policy and uniform within the SR domain.

3. SLID Presence Indicator

The SLID Presence Indicator (SPI) is set by a SLID-capable IPv6 source node to inform transit routers that a SLID is encoded in the packet.

The SPI is encoded as a specific bit or range of values in the Traffic Class field of the IPv6 header.

The encoding of the SPI in the IPv6 header is governed by local policy and uniform within the SR domain.

4. Ingress PE SLID Assignment

When an ingress PE receives a packet that traverses the SR domain, it encapsulates the packet in an outer IPv6 header and optional SRH as defined in [[RFC8754](#)]. The ingress PE MAY also classify the packet into a slice and set the slice identifier as follows:

- *Set the SPI in the outer IPv6 header.

- *Write the SLID in the outer IPv6 header.

The slice classification method is outside the scope of this document.

5. Per-Slice Forwarding

Any router within the SR domain that forwards a packet with SPI bit set uses the SLID to select a slice and apply per-slice policies.

There are many different policies that could define a slice for a particular application or service. The most basic of these is bandwidth-allocation, an implementation complying with this specification SHOULD support the bandwidth-allocation slice as defined in the next section.

6. Bandwidth-Allocation Slice

A per-slice policy is configured at each interface of each router in the SR domain, with one traffic shaper per SLID. The bitrate of each shaper is configured to reflect the bandwidth allocation of the per-slice policy.

If shapers are not available, or desirable, an implementation MAY configure one scheduling queue per SLID with a guaranteed bandwidth equal to the bandwidth-allocation for the slice. This option allows a slice to consume more bandwidth than its allocation when available.

Per-slice shapers or queues effectively provides a virtual port per slice. This solution MAY be complemented with a per-virtual-port hierarchical DiffServ policy. Within the context of one specific slice, packets are further classified into children DiffServ queues which hang from the virtual port. The DSCP value in the IPv6 header SHOULD be used for queue selection.

7. Backward Compatibility

The Flow Label usage described in this document is consistent with [\[RFC6437\]](#) and [\[RFC6438\]](#).

PE routers that do not set the SPI do not enable the SLID semantic of the Flow Label bits. Hence, SLID-aware routers would not attempt to classify these packets into a slice.

Any router that does not process the SPI nor the SLID forwards packets as usual.

8. Acknowledgements

The authors would like to thank Darren Dukes, Ketan Talaulikar, Jisu Bhattacharya, John Bettink, Aman Manot, and David Melman for their insightful feedback on this document.

9. References

9.1. Normative References

- [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>>.
- [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>>.

9.2. Informative References

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

Saad, T., Beeram, V. P., Wen, B., Ceccarelli, D., Halpern, J., Peng, S., Chen, R., Liu, X., Contreras, L. M., and R. Rokui, "Realizing Network Slices in IP/MPLS Networks", Work in Progress, Internet-Draft, draft-bestbar-teas-ns-packet-03, 11 July 2021, <<https://www.ietf.org/archive/id/draft-bestbar-teas-ns-packet-03.txt>>.

[I-D.ietf-teas-ietf-network-slices]

Farrel, A., Gray, E., 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-03, 23 May 2021, <<https://www.ietf.org/archive/id/draft-ietf-teas-ietf-network-slices-03.txt>>.

[RFC6437] Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <<https://www.rfc-editor.org/info/rfc6437>>.

[RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <<https://www.rfc-editor.org/info/rfc6438>>.

Authors' Addresses

Clarence Filsfils (editor)
Cisco Systems, Inc.
Belgium

Email: cf@cisco.com

Francois Clad (editor)
Cisco Systems, Inc.
France

Email: fclad@cisco.com

Pablo Camarillo
Cisco Systems, Inc.
Spain

Email: pcamaril@cisco.com

Kamran Raza
Cisco Systems, Inc.
Canada

Email: skraza@cisco.com

Daniel Voyer
Bell Canada
Canada

Email: daniel.voyer@bell.ca

Reza Rokui
Nokia
Canada

Email: reza.rokui@nokia.com