SPRING Working Group
Internet Draft
Intended status: Standards Track
Expires: September 03, 2025

J. Ye
China Mobile
C. Lin
New H3C Technologies
D. Lu
M. Chen
China Mobile

March 3, 2025

SRv6 Context Indicator SIDs for SR-Aware Services
draft-lin-spring-srv6-aware-context-indicator-05

Abstract

A context indicator provides the context on how to process the packet for service nodes. This document describes how to use SRv6 SIDs as context indicator for SR-aware services. The corresponding Endpoint behaviors are defined.

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 September 03, 2025.

Table of Contents

1. Introduction

Segment Routing (SR) [RFC8402] leverages the source routing paradigm. A node steers a packet through an SR Policy instantiated as an ordered list of instructions called "segments". Segment Routing (SR) can be applied to the IPv6 data plane using Segment Routing Header (SRH) [RFC8754], which is called SRv6.

The segments may encode simple routing instructions for forwarding packets along a specific network path, but also steer them through VNFs or physical service appliances available in the network. [I-D.ietf-spring-sr-service-programming] describes how a service can be associated with a SID (Segment Identifier) and how these service SIDs are integrated within an SR policy.

Services are categorized in two types, SR-aware and SR-unaware services. An SR-aware service can process the SR information in the packets it receives. [I-D.ietf-spring-sr-service-programming] defines an SRv6 Endpoint Behavior, End.AN, for SR-aware function. But service-specific functions are not defined.

A context indicator provides the context on how to process the packet for service nodes. A SID can be associated with such function of context indicator in the SR-aware services. For example, an SR-aware firewall may use a context indicator SID to identify the specific virtual firewall instance when applying VPN-specific rules for inner packets.

In some cases, a context indicator SID can be dynamically associated with a bunch of contexts, and indicate the SR-aware service nodes to identify the particular context with additional information carried in the packet. When such dynamic context indicator SID is contained in the SR Policy, the headend node will fill additional context information in the corresponding field of the packet, based on which traffic flow the packet belongs to. The End.AN SID defined in [I-D.ietf-spring-sr-service-programming] is not adequate for context indicators, especially for the dynamic ones. Because End.AN is too abstract and general for the headend node to determine its actions.

This document describes how to use SRv6 SIDs as context indicator for SR-aware services. These SIDs are called SRv6 Context Indicator SIDs. The corresponding Endpoint behaviors for SRv6 Context Indicator SIDs are defined in this document.

2. Use Case

In traditional security resource pool, Policy-Based Routing (PBR) is employed to orchestrate Service Function Chain(SFC). To differentiate between tenants and subsequently provide them with personalized value-added service, VLANs are deployed in different 3-layer sub-interfaces, which are bound to distinct vpn instances. These VLANs and VPN instances serve the purpose of isolating tenants from one another. Despite the traffic of some tenants needs to be processed by virtual firewall, their specific service requirements may vary. For instance, as in shown in figure 1, the traffic of tenant A enters vFW through a 3-layer sub-interface with VLAN ’a’, while tenant B’ s traffic enters through that of VLAN ’b’. With the existence of VLANs and VPN instances, vFW would be able to figure out tenants and select their exclusive vsys instances to supply service. For example, the traffic of tenant A with destined for address X would be dropped by vsys1 instance on vFW, while tenant B’s traffic with destination address Y will be dropped by vsys2 instance. This PBR method necessitates manual configuration and has drawbacks such as complexity in configuration.

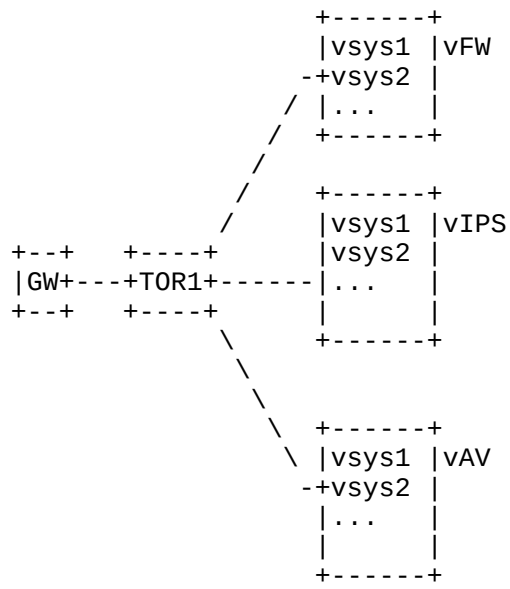

Figure 1

With the emergence of SRv6, its inherent programmability makes it suitable for SFC orchestration. However, the current SRv6 SID could accurately steer packets to a specified service node, for instance, through the utilization of END.AN, but it falls short in communicating to the node the specific service (e.g. vsys instance in security resource pool scenario) that the packets require. Therefore, it is of significant importance to extend and specify END.AN.

3. SRv6 Context Indicator SIDs

An SRv6 Context Indicator SID is associated with a local context on the SR-aware service node. It instructs the node to process the packet by using the specific context.

This document defines new types of Endpoint behaviors for SRv6 Context Indicator SID, End.AN.CI.S and End.AN.CI.D (including End.AN.CI.D.A, End.AN.CI.D.T, End.AN.CI.D.V, and End.AN.CI.D.D), which are variants of the End.AN behavior [I-D.ietf-spring-sr-service-programming]. End.AN.CI.S is statically associated with one particular context. End.AN.CI.D (including End.AN.CI.D.A, End.AN.CI.D.T, End.AN.CI.D.V, and End.AN.CI.D.D) is dynamically associated with a bunch of local contexts, and additional variable

information carried in the packet is used to identify the particular context.

3.1. End.AN.CI.S: SR-Aware Service Static Context Indicator

The "Endpoint with SR-Aware Service Static Context Indicator" behavior ("End.AN.CI.S" for short) is a variant of the End.AN behavior.

One of the applications of the End.AN.CI.S behavior is the SR-aware firewall use case where the associated context identifies a specific virtual firewall instance.

When N receives a packet whose IPv6 DA is S and S is a local End.AN.CI.S SID associated with a local context C, N does the following:

```
S01. When an SRH is processed {
S02.   If (Segments Left == 0) {
S03.      Proceed to process the next header in the packet.
S04.   }
S05.   If (IPv6 Hop Limit <= 1) {
S06.      Send an ICMP Time Exceeded message to the Source Address
          with Code 0 (Hop limit exceeded in transit),
          interrupt packet processing, and discard the packet.
S07.   }
S08.   max_LE = (Hdr Ext Len / 2) - 1
S09.   If ((Last Entry > max_LE) or
           (Segments Left > Last Entry+1)) {
S10.      Send an ICMP Parameter Problem to the Source Address
             with Code 0 (Erroneous header field encountered)
             and Pointer set to the Segments Left field,
             interrupt packet processing, and discard the packet.
S11.   }
S12.   Set the packet’s associated context to C and perform service
S13.   Decrement IPv6 Hop Limit by 1
S14.   Decrement Segments Left by 1
S15.   Update IPv6 DA with Segment List[Segments Left]
S16.   Submit the packet to the egress IPv6 FIB lookup for
          transmission to the new destination
S17. }
```

3.2. End.AN.CI.D: SR-Aware Service Dynamic Context Indicator

The "Endpoint with SR-Aware Service Dynamic Context Indicator" behavior ("End.AN.CI.D" for short) is a variant of the End.AN behavior.

When N receives a packet whose IPv6 DA is S and S is a local End.AN.CI.D SID, the line S12 from the End.AN.CI.S processing is replaced by the following:

```
S12.   Set the packet’s associated context by using variable
          context information carried in the packet and
          perform service.
S13.   If (the context information cannot be understood) {
S14.      Send an ICMP Parameter Problem to the Source Address
             with Code 0 (Erroneous header field encountered)
             and Pointer set to the context information field,
             interrupt packet processing, and discard the packet.
S15.   }
```

There are four sub-types of End.AN.CI.D SID, carrying variable context information associated with the End.AN.CI.D SID in different positions:

- End.AN.CI.D.A: Arguments in SID
- End.AN.CI.D.T: SRH Tag
- End.AN.CI.D.V: SRH TLV for context
- End.AN.CI.D.D: New options in DoH before SRH

3.2.1. End.AN.CI.D.A: SR-Aware Service Dynamic Context Indicator with Variable Context Information in Arguments

The behavior also takes an argument: "Arg.VCI". This argument provides variable context information for service. In this case, the line S12 from the End.AN.CI.D processing is as the following:

```
S12.   Set the packet’s associated context by using variable
          context information carried in the Arg.VCI and
          perform service
```

3.2.2. End.AN.CI.D.T: SR-Aware Service Dynamic Context Indicator with Variable Context Information in SRH Tag

The Tag field in SRH could be used to carry variable context information. In this case, the line S12 from the End.AN.CI.D processing is as the following:

S12. Set the packet’s associated context by using variable context information carried in the SRH Tag and perform service

3.2.3. End.AN.CI.D.V: SR-Aware Service Dynamic Context Indicator with Variable Context Information in SRH TLV

Optional TLV in SRH could be extended for variable context information, which is used together with End.AN.CI.D. The Context Information TLV has the following format:

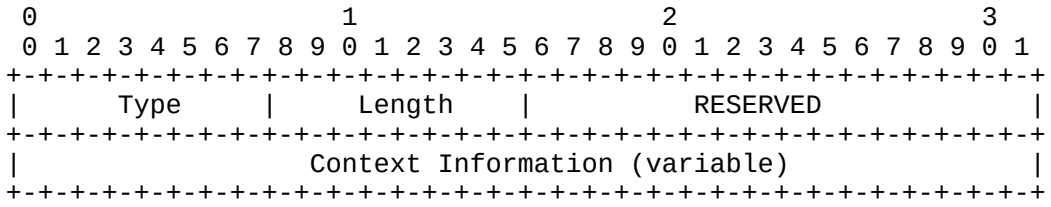

In this case, the line S12 from the End.AN.CI.D processing is as the following:

S12. Set the packet’s associated context by using variable context information carried in the SRH Context Information TLV and perform service

3.2.4. End.AN.CI.D.D: SR-Aware Service Dynamic Context Indicator with Variable Context Information in DOH

Variable context information could also be carried through DOH for the specified segment. The definition of such DOH Option is outside the scope of this document.

In this case, the line S12 from the End.AN.CI.D processing is as the following:

S12. Set the packet’s associated context by using variable context information carried in the DOH and perform service

4. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to [RFC7942].

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in

[RFC7942]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC7942], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

4.1. H3C’s Commercial Delivery

The feature has been implemented on H3C Comware V7.

* Organization: H3C
* Implementation: H3C’s Commercial Delivery implementation based on Comware V7.
* Description: The implementation has been done.
* Maturity Level: Product
* Contact: linchangwang.04414@h3c.com

5. Security Considerations

TBD

6. IANA Considerations

This I-D requests the IANA to allocate, within the "SRv6 Endpoint Behaviors" sub-registry belonging to the top-level "Segment-routing with IPv6 dataplane (SRv6) Parameters" registry, the following allocations:

Value	Description	Reference
0x00A9	End.AN.CI.S	[This.ID]
0x00AA	End.AN.CI.D.A	[This.ID]
0x00AB	End.AN.CI.D.T	[This.ID]
0x00AC	End.AN.CI.D.V	[This.ID]
0x00AD	End.AN.CI.D.D	[This.ID]

7. References

7.1. Normative References

[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/rfc/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/rfc/rfc8754>.

[I-D.ietf-spring-sr-service-programming] Clad, F., Xu, X., Filsfils, C., Bernier, D., Li, C., Decraene, B., Ma, S., Yadlapalli, C., Henderickx, W., and S. Salsano, "Service Programming with Segment Routing", Work in Progress, Internet-Draft, draft-ietf-spring-sr-service-programming-08, 21 August 2023, <https://www.ietf.org/archive/id/draft-ietf-spring-sr-service-programming-08.txt>.

7.2. Informative References

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://www.rfc-editor.org/info/rfc7942>.

Authors' Addresses

Jiaming Ye
China Mobile
China
Email: yejiaming@chinamobile.com

Changwang Lin
New H3C Technologies
China
Email: linchangwang.04414@h3c.com

Dongjie Lu
China Mobile
China
Email: ludongjie@chinamobile.com

Meiling Chen
China Mobile
China
Email: chenmeiling@chinamobile.com