

Workgroup: INTAREA
Internet-Draft:
draft-computing-segment-for-service-routing-00
Published: 10 October 2022
Intended Status: Standards Track
Expires: 13 April 2023
Authors: F. Zhou D. Yuan
 ZTE Corporation ZTE Corporation
 Computing Segment for Service Routing in SAN

Abstract

Since services delivered from cloud need delicate coordination among the client, network and cloud, this draft defines a new Segment to provide service routing and addressing functions by leveraging SRv6 Segment programming capabilities. With Computing Segments proposed, the network gains its capability to identify and process SAN header in need and a complete service routing procedure can be achieved.

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 13 April 2023.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in

Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
 - [1.1. Service Identification in SAN](#)
 - [1.2. Service Routing in SAN](#)
- [2. Requirements Language](#)
- [3. Terminology](#)
- [4. Computing Segment](#)
 - [4.1. When a SAN Header is Carried as an Option in the HBH](#)
 - [4.2. When a SAN Header is Carried as an Option in the DOH](#)
 - [4.3. When a SAN Header is Carried as a Type of SRH TLV](#)
- [5. Use Case](#)
- [6. Security Considerations](#)
- [7. Acknowledgements](#)
- [8. IANA Considerations](#)
- [9. Normative References](#)
- [Authors' Addresses](#)

1. Introduction

1.1. Service Identification in SAN

In order to deliver responsive services to clients, computing resources continuously migrate and spread from central sites to edge nodes. As shown in Figure 1, multiple instances located distributedly in different resource pools are capable of providing services. Compared with applying traditional IP routing protocols, a fine-grained service routing policy is capable of achieving optimal and efficient invocation of both computing power and the network.

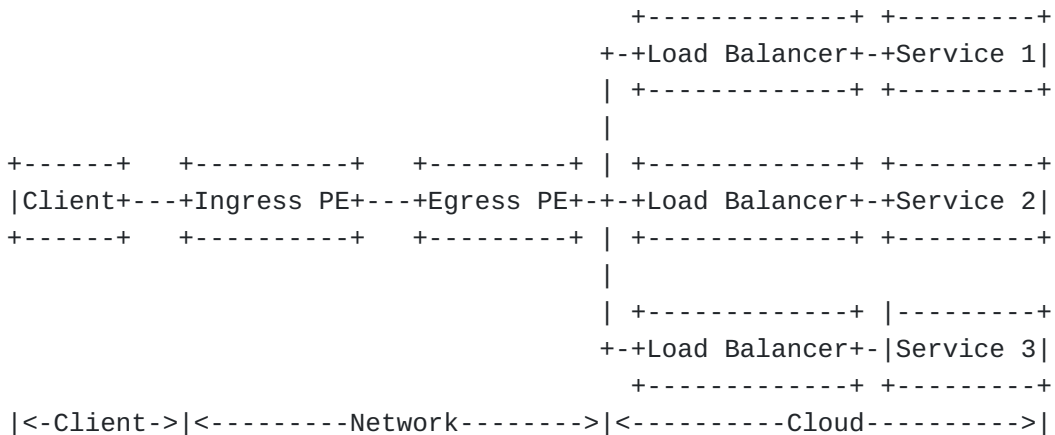


Figure 1: Computing Power Networks

In order to implement service routing, the network should be aware of specific services and a service awareness network framework is introduced in [[I-D.huang-service-aware-network-framework](#)]. Within the proposed network framework, a service identification is defined as a SAN ID(Service ID) in [[I-D.service-identification-header-of-san](#)] to represent a globally unique service semantic identification.

As mentioned in [[I-D.encapsulation-of-san-header](#)], a SAN ID is encapsulated in a SAN header which can be carried as an option in the IPv6 Hop-by-Hop Options Header, Destination Options Header and a type of SRH TLV. Since services delivered from cloud need delicate coordination among the client, network and cloud and thus simply encapsulating SAN header among packets delivery can hardly satisfy various practical situations:

- *The Destination Options header is used to carry optional information that need be examined by the destination of the path which is defined in [[RFC8200](#)], SAN header will only be resolved by the destination node. When a multi-layer routing protocol is applied in the network domain, a quantity of relay nodes besides the destination are required to identify SAN ID and forward the received packet accordingly as well. Thus, simply carrying a SAN header can not fulfill a multi-layer service routing procedure.
- *When a SAN header is carried as an option in the IPv6 Hop-by-Hop Options Header, it may be processed by each nodes. Practically, not all nodes along the delivery path of the packet are capable of identifying and processing a SAN header. The SAN header may be modified and changed and the packet may even be discarded in the forwarding process.
- *The Segment Routing Header (SRH) and the SRH TLV is defined in [[RFC8754](#)]. Since the segment list is encoded in order, it must be orchestrated in advance which indicates various endpoint behaviours in order to successively implement the designated service routing. Previous orchestration should be regarded to be severe restrictions.

To achieve a SAN header being processed in need in the network domain and to preserve its identifiability along the path from the client to the server, a new Segment to specify and standardize node behaviours is urgently required.

1.2. Service Routing in SAN

As shown in Figure 2, a service routing table is designed to establish a mapping relationship between the SAN ID and the conventional IP routing table.

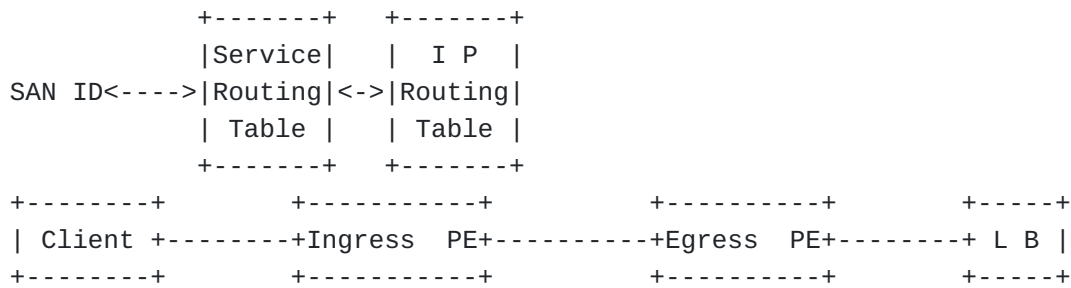


Figure 2: Service Routing in SAN

A service routing table can be published from a control and management system to the network domain within a centralized control plane while it can also be calculated and generated by the Ingress PE itself under a distributed control plane.

With considerations of both path metrics and service SLA requirements, a specific service routing table is introduced, including multiple attributes, SAN ID and outer gateway for instance. Afterwards, a corresponding IP routing table should be indexed which further determines the next hop or an SRv6 policy.

In order to describe and standardize the mentioned behaviours, a new Computing Segment is proposed. With Computing Segments, multiple nodes in the network domain can be informed to locate and identify SAN header in need and to implement a referred forwarding behaviour through the complete procedure.

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

*SAN: Service Aware Network

*SAN ID: Service Aware Network Identification, an identification designed to indicate the fundamental and common service types

*SAN header: Encapsulation format of the SAN ID

*DOH: Destination Options Header

*HBH: Hop-by-Hop Options Header

*SRH: Segment Routing Header

*SID: Segment Identifier

*FIB: Forwarding Information Base

*DA: Destination Address

*LB: Load Balancer

4. Computing Segment

This draft introduces a new SRv6 Segment, namely Computing Segment, aiming to describe the behaviour of querying service routing table and corresponding packet forwarding.

Computing Segment is the identifier of packets in which a corresponding SAN header should be identified and further being forwarded via the matched service routing table entity, indicating the following operations:

- *Identify the SAN ID encapsulated in DOH, HBH or SRH TLV.
- *Query the forwarding table entry indexed by SAN ID.
- *Forward the packet to the new destination.

In the case of SRv6, a new behavior End.C for Computing Segment is defined. An instance of a Computing SID is associated with a service routing table and a source address.

Behaviours of End.C when a SAN header is carried as an option in the HBH, DOH or a type of SRH TLV are described in the following sections.

4.1. When a SAN Header is Carried as an Option in the HBH

When an IPv6 node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as a SID (S), and S is a Computing SID, N does:

```

S01.  When an IPv6 packet is processed {
S02.    Identify the SAN ID encapsulated in the option of the HBH
S03.    Query the forwarding table entry indexed by SAN ID
S04.    Set the packet's associated FIB table to the specific FIB
S05.    Set the IPv6 DA to the next hop
S06.    Maintain the TLVs in the HBH
S07.    Resubmit the packet and transmit to the new destination
S08.  }

```

Figure 3: When a SAN Header is Carried as an Option in the HBH

4.2. When a SAN Header is Carried as an Option in the DOH

When an IPv6 node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as a SID (S), and S is a Computing SID, N does:

```

S01.  When an IPv6 packet is processed {
S02.    Identify the SAN ID encapsulated in the option of the DOH
S03.    Query the forwarding table entry indexed by SAN ID
S04.    Set the packet's associated FIB table to the specific FIB
S05.    Set the IPv6 DA to the next hop
S06.    Maintain the TLVs in the DOH
S07.    Resubmit the packet and transmit to the new destination
S08.  }

```

Figure 4: When a SAN Header is Carried as an Option in the DOH

4.3. When a SAN Header is Carried as a Type of SRH TLV

When an SRv6-capable node (N) receives an IPv6 packet whose destination address matches a local IPv6 address instantiated as an SRv6 SID (S), and S is a Computing SID, N does:

```

S01.  When an SRH is processed {
S02.      If (Segments Left>0) {
S03.          Decrement IPv6 Hop Limit by 1
S04.          Decrement Segments Left by 1
S05.          Update IPv6 DA with Segment List[Segments Left]
S06.          Identify the SAN ID encapsulated in the SRH TLV
S07.          Query the forwarding table entry indexed by SAN ID
S08.          Set the packet's associated FIB table to the specific
S09.          Maintain the TLVs in the SRH
S10.          Resubmit the packet transmit to the new destination
S11.      }
S12. }

```

Figure 5: When a SAN Header is Carried as a Type of SRH TLV

When a SAN header is carried as a type of SRH TLV, Computing SIDs in Segment List are required to be orchestrated in advance which previously indicates the the determinism of a multi-segment routing policy. Therefore, Computing Segment does not cooperate well with the circumstances when a SAN header is carried as a type of SRH TLV.

5. Use Case

When a SAN header is carried as an option in the DOH, a typical service addressing procedure is shown in Figure 6.

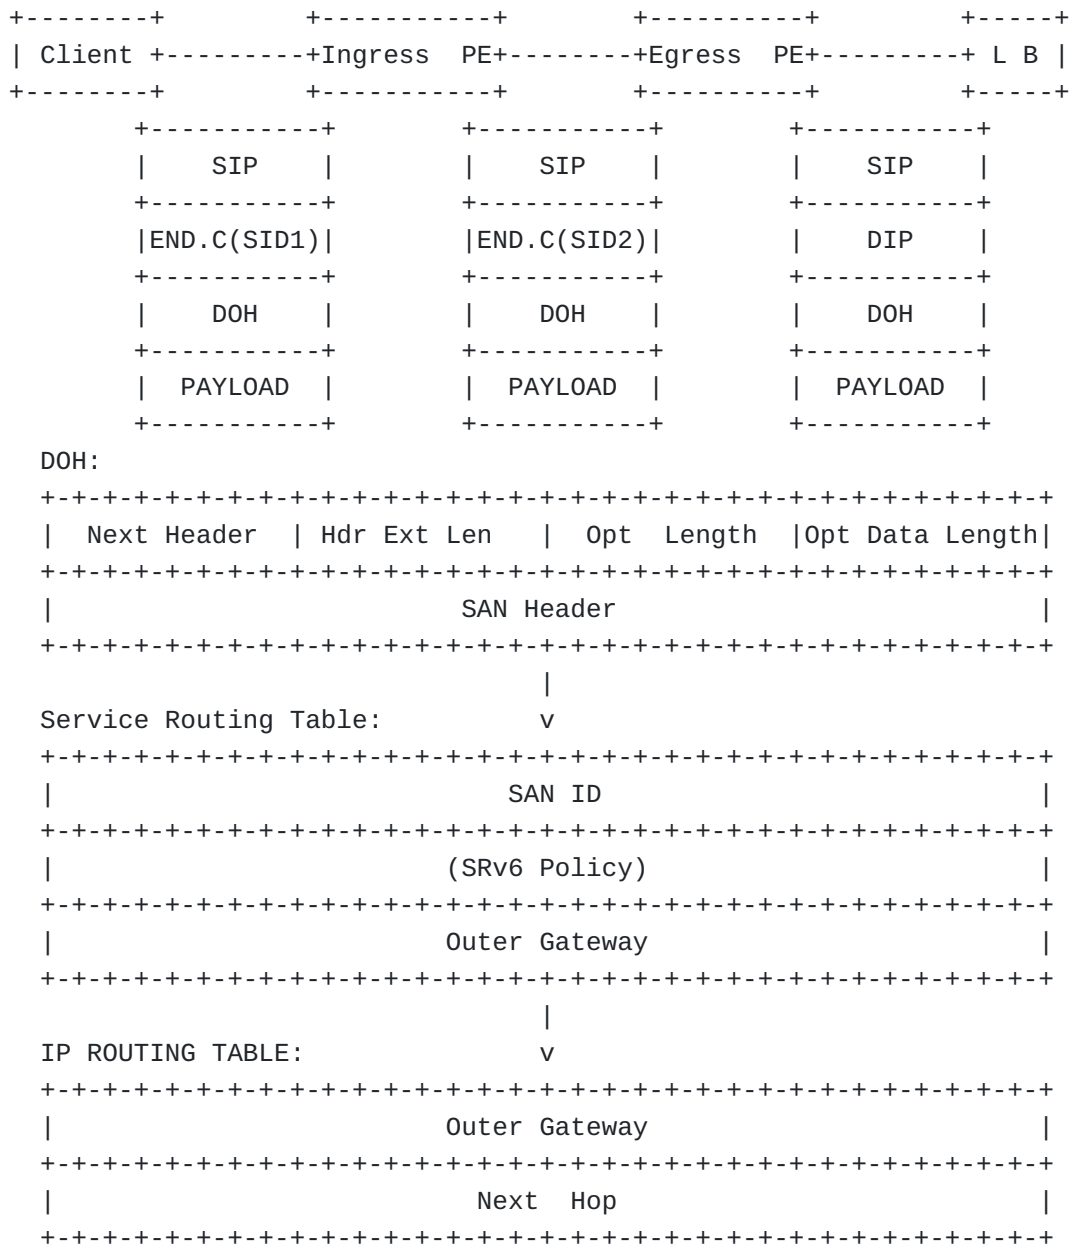


Figure 6: Typical Service Addressing Procedure with Service ID Encapsulated in the DOH

Suppose the Endpoint behaviour of END.C is configured at Ingress PE and Egress PE, namely SID 1 and SID 2 respectively. SID1 and SID2 are advertised to the nodes in the network by IGP. The service addressing procedure from the client to the cloud is described below:

The Computing SID of Ingress PE (SID1) is configured as DA by the client. The packet carrying the SAN header as the option of the DOH is forwarded to Ingress PE.

When Ingress PE receives the packet, it queries the local routing table in accordance with DA and identifies that DA is a Computing SID (SID1). As defined in 4.2, the Ingress PE continues to forward the packet carrying the DOH.

When Egress PE receives the packet, it queries the local routing table in accordance with DA and identifies that DA is a Computing SID (SID2). As defined in 4.2, the Egress PE continues to forward the packet carrying the DOH.

When an intra-cloud LB receives the packet, the packet can be forwarded in accordance with the Endpoint behaviour defined in 4.2. or be processed as a normal IPV6 packet, depending on the practical circumstances.

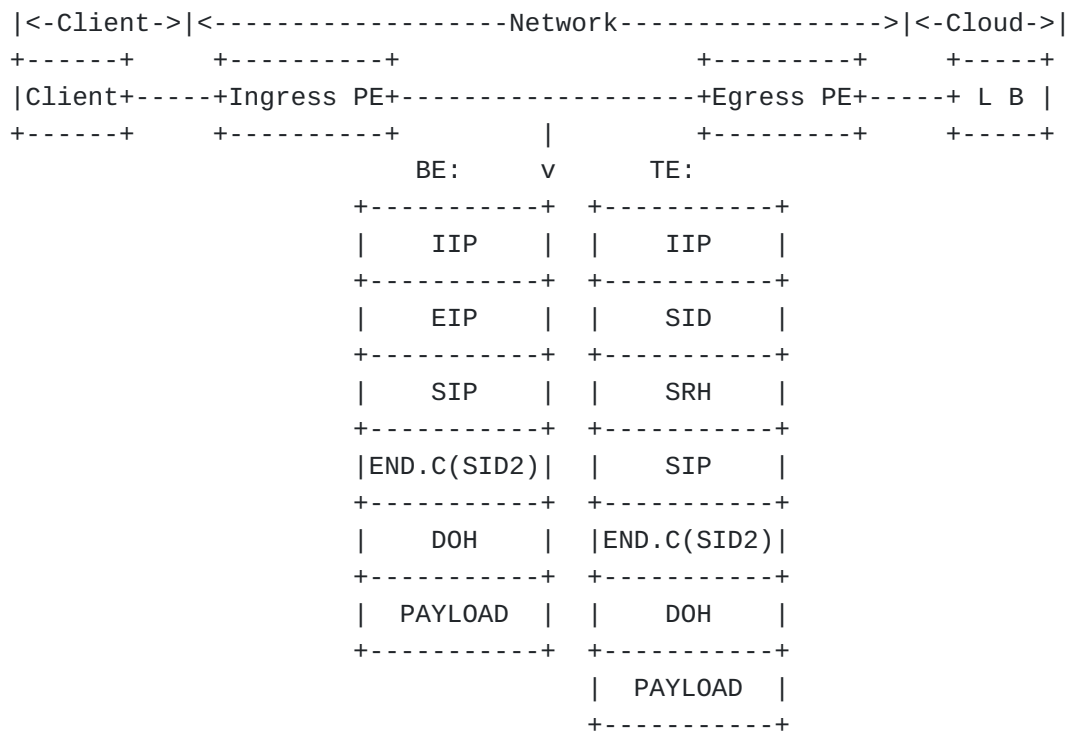


Figure 7: Outer Headers Encapsulated between Ingress PE and Egress PE

As shown in Figure 7, between Ingress PE and Egress PE, an outer header including SRH should be encapsulated when the traffic follows a specific SRv6 TE policy. Otherwise, a normal IPv6 header should be encapsulated under a BE condition. In the introduced case, the SAN header is not perceived by relay devices between Ingress PE and Egress PE.

6. Security Considerations

TBA

7. Acknowledgements

TBA

8. IANA Considerations

This document requires registration of End.C behavior in "SRv6 Endpoint Behaviors" sub-registry of "Segment Routing Parameters" registry.

9. Normative References

[I-D.encapsulation-of-san-header] Ma, L., Zhao, D., and F. Zhou, "Encapsulation of SAN Header", Work in Progress, Internet-Draft, draft-encapsulation-of-san-header-00, 18 August 2022, <<https://www.ietf.org/archive/id/draft-encapsulation-of-san-header-00.txt>>.

[I-D.huang-service-aware-network-framework] Huang, D. and B. Tan, "Service Aware Network Framework", Work in Progress, Internet-Draft, draft-huang-service-aware-network-framework-00, 24 May 2022, <<https://www.ietf.org/archive/id/draft-huang-service-aware-network-framework-00.txt>>.

[I-D.service-identification-header-of-san] Ma, L., Zhou, F., and H. Li, "Service Identification Header of Service Aware Network", Work in Progress, Internet-Draft, draft-service-identification-header-of-san-00, 18 August 2022, <<https://www.ietf.org/archive/id/draft-service-identification-header-of-san-00.txt>>.

[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>>.

[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>>.

[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>>.

Authors' Addresses

Fenlin Zhou
ZTE Corporation
No.50 Software Avenue
Nanjing
Jiangsu, 210012
China

Email: zhou.fenlin@zte.com.cn

Dongyu Yuan
ZTE Corporation
No.50 Software Avenue
Nanjing
Jiangsu, 210012
China

Email: yuan.dongyu@zte.com.cn