Network Working Group
Internet Draft
Intended status: Informational
Expires: August 28, 2026

Gary Geng
Tencent
Y. Liu
China Mobile
C. Xie
China Telecom
C. Lin
New H3C Technologies

February 26, 2026

Considerations for traffic steering to SRv6
draft-geng-srv6ops-traffic-steering-to-srv6-03

Abstract

This document primarily describes the traffic steering towards SRv6-BE and SRv6-TE respectively, providing an overview of the main traffic steering methods for these two approaches. Furthermore, it discusses the recommended traffic steering methods for various typical scenarios.

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 February 26, 2026.

Copyright Notice

Table of Contents

1. Introduction

The general purpose of traffic steering is to optimize the allocation and transmission of network resources, ensure a balanced distribution of network traffic, improve network performance, reduce congestion, and increase available bandwidth to provide users with a better network experience.

This document initially describes the traffic steering towards SRv6-BE and SRv6-TE respectively, and outlines the main traffic steering methods for these two approaches. Finally, it discusses the recommended traffic steering methods for various typical scenarios.

1.1. Conventions and 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. Steering to SRv6 Options

Steering to SRv6 can be categorized into two types: Steering based on destination address and Steering based on flow characteristics.

The means of traffic steering in SRv6 include using static routing for traffic steering, employing PBR (Policy-Based Routing) policies for traffic steering, distributing routes through BGP (Border Gateway Protocol) for traffic steering, utilizing BGP-Flowspec to publish rules for traffic steering, and utilizing IGP-Shortcut for traffic steering.

2.1. Steering to SRv6 based on destination address

Traffic can typically be steered based on the destination address by matching traffic destination address via static routing or utilizing Policy-Based Routing (PBR).

1)Steering traffic to SRv6 via static routing:

ipv6 route-static {x:x::x:x/xx} {y:y::y:y}

Steering traffic based on the destination address with static routing.

Where {y:y::y:y} represents an SRv6 SID, which can be a DT4/DT6 address, indicating traffic steering into an L3VPN;

If {y:y::y:y} is a BSID address, it represents traffic steering into an SRv6 TE Policy;

ipv6 route-static {x:x::x:x/xx} color {color} end-point ipv6 {end-point}

Through the above static route configuration, for matched destination addresses, traffic color and end-point can be specified, and associated with an SRv6 TE Policy, enabling traffic steering into the SRv6 TE Policy.

2) Steering traffic to SRv6 via PBR:

```
ipv6 policy-based-route srv6 permit node 0

  if-match acl 2000

  apply next-hop y:y::y:y
```

Steering traffic based on the destination address using PBR.

Similarly, if {y:y::y:y} represents a DT4/DT6 address, it indicates traffic steering into an L3VPN;

If {y:y::y:y} is a BSID address, it signifies traffic steering into an SRv6 TE Policy;

When using PBR for traffic steering, for matched destination addresses, specifying traffic color and end-point, and associating with an SRv6 TE Policy, can effectively steer traffic into the SRv6 TE Policy.

3) Steering traffic to SRv6 via BGP-FlowSpec:

By deploying BGP-FlowSpec rules from the controller, traffic matching specific destination addresses can be steered.

Once matched, BGP-FlowSpec can specify the next-hop address as a DT4/DT6 address to route the traffic into an L3VPN. Alternatively, specifying the next-hop address as a BSID address can direct the traffic into an SRv6 TE Policy.

4) Steering traffic to SRv6 via IGP shortcut:

IGP shortcut, also known as the automatic traffic announcement feature of an IGP, treats an SRv6 TE Policy as a direct link between the endpoints for announcement purposes. During route calculation, if the destination address of the traffic corresponds to the tunnel’s destination address, the traffic is steered into the SRv6 TE Policy.

2.2. Steering to SRv6 based on flow characteristics

Identifying traffic based on specific flow characteristics and steering traffic according to these characteristics. Flow characteristics include Layer 2 attribute 802.1p, Layer 3 IP feature

DSCP value, as well as service-level attributes such as service class and TE class ID.

1) Steering Traffic to SRv6 via PBR Based on Traffic Characteristics

```
ipv6 policy-based-route srv6 permit node 0

 if-match acl 2000

 apply next-hop y:y::y:y

#
```

By specifying traffic characteristics in the ACL to match traffic and then designating the next-hop for the traffic as the SRv6 next-hop address or BSID address, traffic can be directed to SRv6 BE or SRv6 TE policy.

2)Steering traffic via BGP flowspec

By using flowspec, specify the next-hop address in the route attributes as the BSID of the SRv6 TE Policy, in order to steering the traffic associated with this route to the destination of the SRv6 TE Policy.

By using flowspec, specify the next-hop address in the route attributes as the BSID of the SRv6 TE Policy, in order to steering the traffic associated with this route to the destination of the SRv6 TE Policy.

3) Steering to SRv6 TE via Tunnel-Policy

First, based on the traffic destination address, tunnel policies are matched and associated with an SRv6 TE Policy-Group. The Policy-Group contains multiple policies, each with a different color value. A specific traffic characteristic is then mapped to different color values, and based on the color value, corresponding SR Policies are found within the Policy-Group to direct the traffic to the SR TE Policy. Typically, the characteristics used to map traffic to different colors include DSCP values, Dot1P values, TE Class IDs, service classes, and others.

3. Considerations

SRv6 BE(Best Effort), is a method of forwarding traffic without strict quality of service guarantees. It allows for the flexible and efficient forwarding of packets, prioritizing simplicity and scalability. This approach is well-suited for scenarios where fine-grained traffic control is not necessary and where best effort delivery meets the requirements of the network.

Traffic engineering technology calculates and arranges the forwarding paths of traffic to optimize network resource utilization and improve bandwidth efficiency. Additionally, this technology ensures reliable service quality for business operations and prevents all business traffic from competing for resources on the shortest path. Therefore, deploying and utilizing traffic engineering in SRv6 networks has become a necessary requirement for the promotion and development of SRv6 technology.

The traffic engineering technology based on SRv6 is referred to as SRv6 TE.

As shown in Figure 1, SRv6 Policy 1 has a BSID of 1000::1, Color 100, and Endpoint 4::4, with a forwarding path of B->C->D. SRv6 Policy 2 has a BSID of 2000::1, Color 200, and Endpoint 4::4, with a forwarding path of E->F->D.

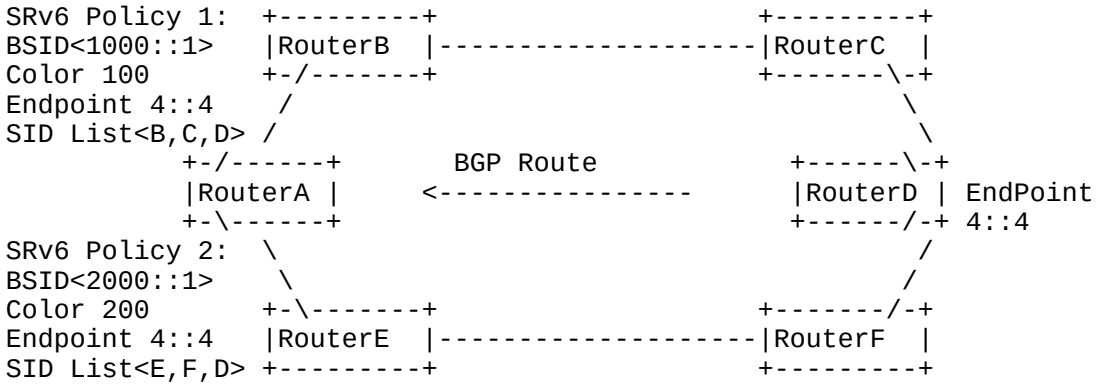

Figure 1. SRv6 Traffic Steering Network Diagram

3.1. Traffic steering based on destination address

3.1.1. BSID-based Traffic Steering

When a device receives a packet with a destination IPv6 address matching the BSID of an SRv6 TE Policy, the packet will be forwarded according to the corresponding SRv6 TE Policy. BSID-based traffic steering is commonly used in BSID stitching scenarios, where the BSID of another SRv6 TE Policy is added to the SID list of one SRv6 TE Policy. This helps reduce the length of the SRH header in the packet during the forwarding process, enabling seamless stitching between different SRv6 TE Policies or between an SRv6 TE Policy and an SR-MPLS TE Policy.

As shown in Figure 1, when the destination address of traffic is specified as BSID 1000::1 of SRv6 Policy 1, the traffic will be forwarded along the path defined by SRv6 TE Policy 1. If the destination address is specified as BSID 2000::1 of SRv6 Policy 2, the traffic will be forwarded along the path defined by SRv6 TE Policy 2.

3.1.2. Color-based Traffic Steering

Color-based traffic steering is one of the fundamental methods used in SRv6 TE Policy. This approach leverages the BGP route’s extended community attribute called Color and the destination address to match the Color and End-point address in the SRv6 TE Policy. Typically, if there is an SRv6 TE Policy on the device with the same Color and End-point address as the Color extended community attribute and next-hop address of the BGP route, the BGP route will be steered to that SRv6 TE Policy. When the device receives a packet that matches the BGP route, it is forwarded through the SRv6 TE Policy.

As shown in Figure 1, the BGP protocol advertises the prefix routes that require traffic steering, such as specifying the Color attribute of route 1::1/128 as 100 and the next-hop attribute as 4::4, and specifying the Color attribute of route 2::2/128 as 200 and the next-hop attribute as 4::4. Therefore, for traffic with a destination address of 1::1, it will be forwarded along the path defined by SRv6 TE Policy 1, while for traffic with a destination address of 2::2, it will be forwarded along the path defined by SRv6 TE Policy 2.

3.1.3. IGP-Shortcut Traffic Steering

When using IGP-shortcut for traffic steering based on the destination address, the routing table information for traffic steering is no longer published by the BGP protocol. Instead, it is automatically generated on the head-end device A based on the SRv6 TE Policy. In the scenario illustrated in Figure 1, enabling the IGP-shortcut feature causes the head-end node to automatically generate a route for 4::4/128 and set the egress interface to point to the SRv6 TE Policy when the SRv6 TE Policy status is Up. In the event that the SRv6 TE Policy status changes to Down, the automatically generated route pointing to the SRv6 TE Policy is withdrawn and replaced with forwarding based on SRv6 BE.

3.2. Traffic steering based on flow characteristics

3.2.1. Dscp-based Traffic Steering

The basic principle of DSCP-based traffic steering is to route the packets to the corresponding SRv6 TE Policy based on the DSCP (Differentiated Services Code Point) value of the packet. This traffic steering method requires the deployment of SRv6 TE Policy groups [I-D.ietf-spring-sr-policy-group] and the redirection of traffic to these SRv6 TE Policy groups. After that, specific DSCP values are mapped to corresponding SRv6 TE Policy groups to steer the packets to the desired SRv6 TE Policy.

As shown in Figure 1, the BGP protocol advertises the prefix routes that require traffic steering, specifying the next-hop attribute of route 1::1/128 as 4::4 but not specifying the Color attribute. On device A, an SRv6 TE Policy group is created with an End-point address of device D’s address 4::4. Within the SRv6 TE Policy group, a Color and DSCP mapping relationship is defined, where DSCP 10 maps to Color 100 and DSCP 20 maps to Color 200. Subsequently, a tunnel policy is configured on source node A, binding the SRv6 TE Policy group with the destination address 2.2.2.2. This arrangement ensures that for traffic with a destination address of 1::1, if the DSCP value is 10, it will be forwarded along the path defined by SRv6 TE Policy 1; if the DSCP value is 20, it will be forwarded along the path defined by SRv6 TE Policy 2.

3.2.2. 802.1p-based Traffic Steering

The basic principle of 802.1p-based traffic steering is to route the packets to the corresponding SRv6 TE Policy based on the 802.1p (Priority) value of the packet. This traffic steering method requires the prior deployment of SRv6 TE Policy groups and the

redirection of traffic to these SRv6 TE Policy groups. Subsequently, based on the mapping rules within the SRv6 TE Policy groups, packets with specific 802.1p values are steered to the corresponding SRv6 TE Policy.

As shown in Figure 1, the BGP protocol advertises the prefix routes that require traffic steering, specifying the next-hop attribute of route 1::1/128 as 4::4 but not specifying the Color attribute. On device A, an SRv6 TE Policy group is created with an End-point address of device D’s address 4::4. Within the SRv6 TE Policy group, a Color and 802.1p mapping relationship is defined, where 802.1p 10 maps to Color 100 and 802.1p 20 maps to Color 200. Subsequently, a tunnel policy is configured on source node A, binding the SRv6 TE Policy group with the destination address 2.2.2.2. This setup ensures that for traffic with a destination address of 1::1, if the 802.1p value is 10, it will be forwarded along the path defined by SRv6 TE Policy 1; if the 802.1p value is 20, it will be forwarded along the path defined by SRv6 TE Policy 2.

3.2.3. Service-class-based Traffic Steering

To ensure that all traffic packets can be steered, even if they do not carry DSCP or Dot1p information, the device introduces a local identification called "service-class" to distinguish different classes of service traffic.

Both service-class-based traffic steering and CBTS-based traffic steering are achieved through the service-class identifier. However, service-class-based traffic steering requires the traffic to first enter the SRv6 TE Policy group, and then, based on the mapping rules within the SRv6 TE Policy group, specific packets with service-class identifiers are redirected to the corresponding SRv6 TE Policy.

As shown in Figure 1, the BGP protocol advertises the prefix routes that require traffic steering, specifying the next-hop attribute of route 1::1/128 as 4::4 but not specifying the Color attribute. On device A, an SRv6 TE Policy group is created with an End-point address of device D’s address 4::4. Within the SRv6 TE Policy group, a Color and service-class mapping relationship is defined, where service-class 1 maps to Color 100 and service-class 2 maps to Color 200. Subsequently, a tunnel policy is configured on source node A, binding the SRv6 TE Policy group with the destination address 2.2.2.2. This setup ensures that for traffic with a destination address of 1::1, if the service-class value is 1, it will be forwarded along the path defined by SRv6 TE Policy 1; if the service-class value is 2, it will be forwarded along the path defined by SRv6 TE Policy 2.

3.2.4. TE-class-based Traffic Steering

Due to the limited length of the local identifier "service-class", with the maximum supported value for most devices being 15 and typically not exceeding 127, it cannot effectively differentiate a vast number of services. Therefore, H3C has introduced the use of "TE class ID" as another local identifier. The TE class ID can have a maximum value of 65535, supporting a more diverse range of traffic types.

The basic principle of TE class ID-based traffic steering is to route the packets to the corresponding SRv6 TE Policy based on the TE class ID identifier. This traffic steering method requires the prior deployment of SRv6 TE Policy groups and the redirection of traffic to these SRv6 TE Policy groups. Then, based on the mapping rules within the SRv6 TE Policy groups, packets marked with specific TE class ID values are steered to the corresponding SRv6 TE Policy.

As shown in Figure 1, the BGP protocol advertises the prefix routes that require traffic steering, specifying the next-hop attribute of route 1::1/128 as 4::4 but not specifying the Color attribute. On device A, an SRv6 TE Policy group is created with an End-point address of device D’s address 4::4. Within the SRv6 TE Policy group, a te-class and service-class mapping relationship is defined, where te-class 1 maps to Color 100 and te-class 2 maps to Color 200. Subsequently, a tunnel policy is configured on source node A, binding the SRv6 TE Policy group with the destination address 2.2.2.2. This setup ensures that for traffic with a destination address of 1::1, if the te-class value is 1, it will be forwarded along the path defined by SRv6 TE Policy 1; if the te-class value is 2, it will be forwarded along the path defined by SRv6 TE Policy 2.

3.2.5. Traffic steering via BGP-FlowSpec

In the scenario illustrated in Figure 2, where a controller is present in the network, fine-grained traffic scheduling can be achieved by the controller distributing traffic steering rules. The controller utilizes BGP Flow-Spec (FS) to distribute rules to the head-end node. These rules can be matched based on the destination address or flow characteristics. By setting the action after matching, the controller can specify the Color attribute and next-hop address for the traffic, thereby achieving traffic steering. BGP FlowSpec enables global traffic scheduling with flexibility.

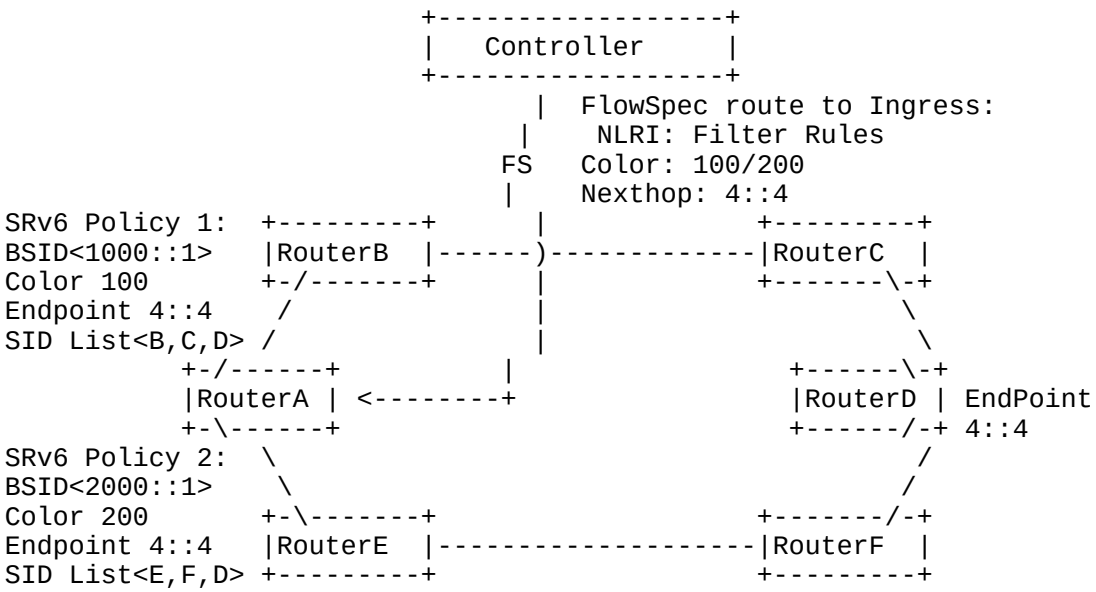

Figure . Steering by BGP FlowSpec

4. Security Considerations

TBD.

5. IANA Considerations

This document makes no request of IANA.

6. References

6.1. 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>.

[RFC9256] C. Filsfils, "Segment Routing Policy Architecture ", July 2022, https://datatracker.ietf.org/doc/rfc9256

6.2. Informative References

[I-D.ietf-spring-sr-policy-group] Cheng, W., Wenying, J., Lin, C., Chen, R., and Y. Zhang, "SR Policy Group", Work in Progress, Internet-Draft, draft-ietf-spring-sr-policy-group-00, 13 January 2026,<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-policy-group-00>.

Authors' Addresses

Gary Geng
Tencent
China

Email: garygeng@tencent.com

Yisong Liu
China Mobile
China

Email: liuyisong@chinamobile.com

Chongfeng Xie
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China

Email: xiechf@chinatelecom.cn

Changwang Lin
New H3C Technologies
Beijing
102209
China

Email: linchangwang.04414@h3c.com