

srv6ops
Internet-Draft
Intended status: Informational
Expires: 18 April 2026

F. Yang
China Mobile
C. Lin
New H3C Technologies
15 October 2025

SRv6 Policy Selector
draft-yang-srv6ops-policy-selector-00

Abstract

Segment routing (SR) [RFC8402] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR Policy is associated with one or more candidate paths, and each candidate path is either dynamic, explicit or composite. This document describes a policy selection mechanism among the candidate SRv6 Policies based on network quality in IPv6 environments.

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 18 April 2026.

Copyright Notice

Copyright (c) 2025 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	2
2. Requirements Language	3
3. Terminology	3
4. Problem and Requirements	3
5. SRv6 Policy Selector	5
5.1. Processing Model	5
5.2. Flow Classification	6
5.3. Flow Steering	6
5.4. Policy Selector	6
5.5. Network Quality Measurement	8
5.6. Flow Forwarding	8
6. Examples of Policy Selector	8
7. IANA Considerations	10
8. Security Considerations	10
9. References	10
9.1. Normative References	10
9.2. Informative References	11
Acknowledgements	11
Authors' Addresses	11

1. Introduction

Segment routing (SR) [RFC8402] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR Policy is associated with one or more candidate paths, and each candidate path is either dynamic, explicit or composite.

The [I-D.ietf-idr-performance-routing] specification defines a mechanism for disseminating path delay information across multiple Autonomous Systems (ASes). This information is used for BGP route computation.

An SRv6 Policy is associated with one or more candidate paths. A composite candidate path acts as a container for grouping SRv6 Policies. As described in section 2.2 in [RFC9256], the composite candidate path construct enables combination of SRv6 Policies, each with explicit candidate paths and/or dynamic candidate paths with

potentially different optimization objectives and constraints, for load-balanced steering of packet flows over its constituent SRv6 Policies. For convenience, the composite candidate path formed by the combination of SRv6 Policies is called parent SRv6 Policy in [I-D.cheng-spring-sr-policy-group].

Different enterprise applications have varying network performance requirements. For instance, voice is highly sensitive to packet loss and jitter, while OA applications are not highly demanding in terms of latency and packet loss.

This document describes a policy selection mechanism among the candidate SRv6 Policies based on network quality in IPv6 environments.

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

The definitions of the basic terms are identical to those found in Segment Routing Policy Architecture [RFC9256]. Office Automation: OA, which stands for "Office Automation," is a collaborative office and internal management platform built using network and software technologies. Parent SR Policy: Refer to [I-D.cheng-spring-sr-policy-group]. A Parent SR Policy is the composite candidate path that acts as a container for grouping SR Policies which meet different service optimization objectives and constraints and have the same destination endpoint.

4. Problem and Requirements

Take the networking shown in Figure 1 below as an example to illustrate the current problems.

CE1 and CE2 are the two access endpoints of the IP telecom network. There are many service flows between CE1 and CE2 that have different requirements for forwarding quality. E.g. OA and voice traffic have different SLA requirement, and expected be carried by different SRv6 Policies. Generally, from CE1 to CE2, voice services with low latency requirements are forwarded along the highly reliable path PE1->PE2->CE2. The OA traffic is forwarded along the normal path PE3->P5->P6->PE2->CE2. When failure or degradation happened in voice traffic SRv6 Policy, there should be possible to assure basic communication for voice traffic by using OA bandwidth.

In single SRv6 Policy, there are many mechanism provide failure/degrade protection, such as TILFA, VPN FRR. However, it is not clear how to handle failure or degradation between multiple SRv6 Policies in a parent SR Policy.

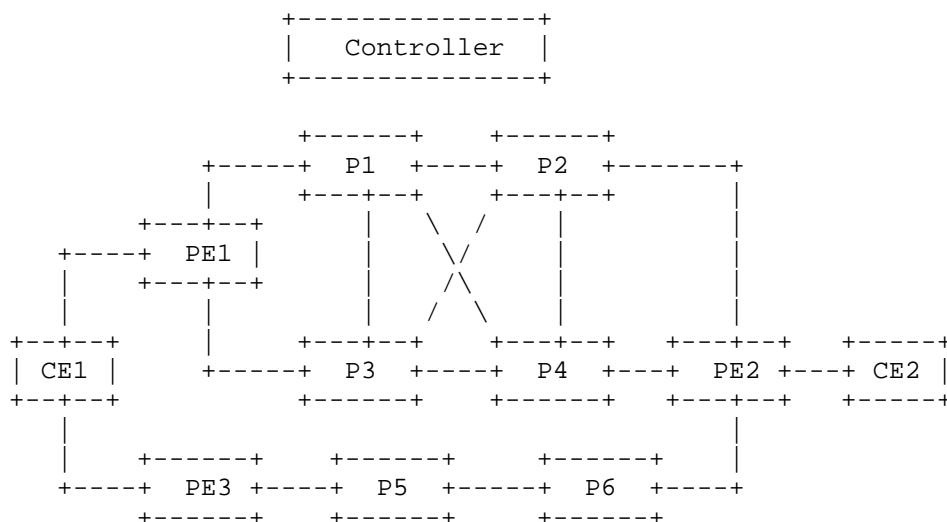


Figure 1

Based on such scenarios, the following requirements are proposed:

1. Maximize failure/degradation protection

In case of failure or degradation detected on one SRv6 Policy, it should be possible to do inter-policy protection.

2. Minimal impact after taking repairing action

Repair action can be done on flow level to minimize the ripple effect cause by forwarding path switchover.

3. Maximize bandwidth efficiency

For some critical applications, it should be possible to forward the traffic over lower class policy in case of higher class SRv6 Policy degradation.

Refer to [I-D.cheng-spring-sr-policy-group], the services with different forwarding quality requirements to the same destination endpoint can be implemented through parent SRv6 Policy group.

This document proposes an SRv6 Policy selector for parent SRv6 Policy based on network quality requirement. The head end node selects the best constituent SRv6 Policy for the application according to the quality of the constituent SRv6 Policy.

Take Figure 1 as an example, there is a parent SRv6 Policy between CE1 to CE2, which has multiple constituent SRv6 Policies. A new SRv6 Policy Selector is defined for the application in the parent SRv6 Policy, which will select best constituent SRv6 Policy in the parent SRv6 Policy. When the head node detects the quality degradation of the active constituent SRv6 Policy, it will select another constituent SRv6 Policy in the parent SRv6 Policy.

5. SRv6 Policy Selector

5.1. Processing Model

A new priority and a new quality threshold is created for each constituent SRv6 Policy. The lower the priority number, the higher the priority. That means active constituent SRv6 Policy will be the one with higher priority and meeting the quality threshold. When the network quality degradation is happened on the active constituent SRv6 Policy, such as the packet loss rate exceeds the threshold, switch to the next high priority constituent SRv6 Policy.

If the quality of the high priority forwarding constituent SRv6 Policy is restored and the specified quality threshold is met, the traffic will be switched back to the high priority constituent SRv6 Policy after a period of wait-to-restore time.

According to the processing logic, the SRv6 Policy Policy Selector model can be divided into five units, including Flow Classification, Flow Steering, Policy Selector, Flow Forwarding, and Network Quality Measurement, as shown in Figure 2 below.




Figure 2

The functions of each unit are described below.

5.2. Flow Classification

After receiving the traffic, the head node first needs to label the traffic with application type according to classification configuration.

5.3. Flow Steering

According to the application type determined by the Flow classification, the header node selects the parent SRv6 Policy for traffic forwarding.

5.4. Policy Selector

SRv6 Policy Selector obtains the current quality of each constituent SRv6 Policy from the Network Quality Measurement unit. Based on the quality threshold and the priority, SRv6 Policy Selector selects the active constituent SRv6 Policy.

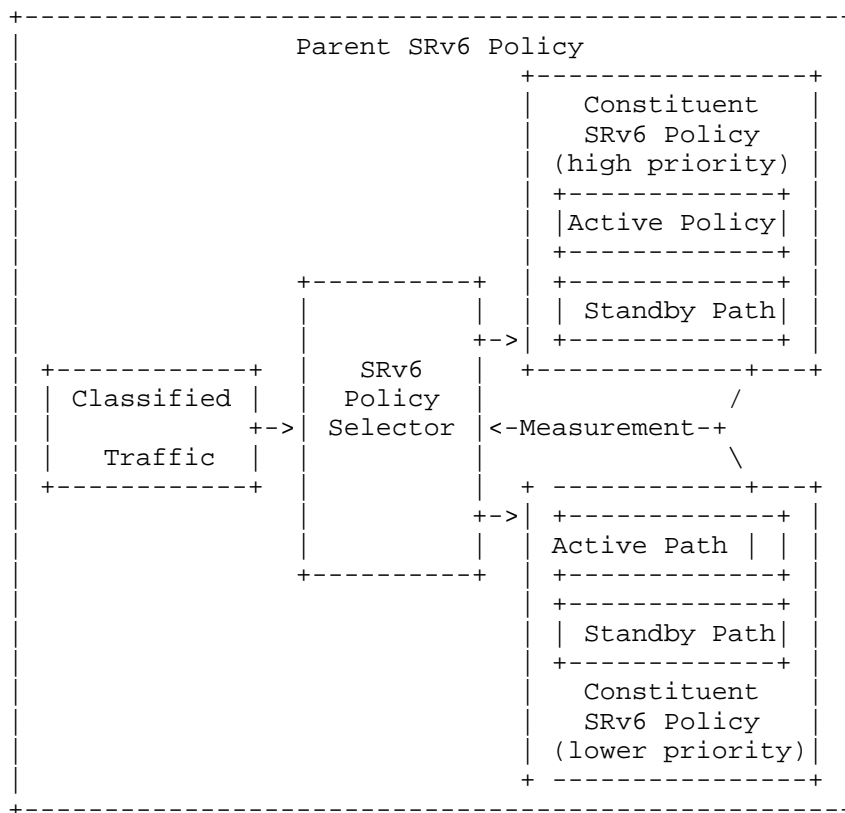


Figure 3

Each parent SRv6 Policy contains multiple constituent SRv6 Policies. Each constituent SRv6 Policy will include two new configuration parameters, "priority" and "threshold". A lower priority value indicates a higher precedence. The constituent SRv6 Policy with the highest priority and qualified threshold will be active.

To avoid frequent path switching when the network quality is unstable, a wait-to-restore timer is required. Only after automatic restore is allowed and the wait-to-restore timer is timeout, the forwarding path switch from the current constituent SRv6 Policy to the one with higher priority.

5.5. Network Quality Measurement

The Network Quality Measurement unit regularly monitors the quality of all effective forwarding paths according to the measurement cycle, records the current performance measurement data of the path, and reports it to the SRv6 Policy Selector unit, which decides whether to switch paths.

The following network quality parameters can be used:

- * Jitter
- * Latency
- * Packet loss
- * Available bandwidth
- * Bandwidth utilization
- * Current traffic statistics
- * Other forwarding performance parameters

The quality parameters can be obtained through active or passive performance measurement methods, such as iOAM, STAMP, TWAMP, SRv6 bandwidth measurement[I-D.liu-ippm-srv6-bandwidth-measurement], etc. The network quality parameters can be calculated by the controller and distributed to the head end node, or calculated by the head end node according to the network measurement data. The measurement method and quality parameter acquisition method are beyond the scope of this document.

5.6. Flow Forwarding

The service flow is forwarded according to the path determined by the Policy Selector unit.

When there are multiple paths with the same priority, the traffic will share the load among these SRv6 Policy paths with the same priority according to the weight value.

6. Examples of Policy Selector

The application of Policy Selector is described in detail in L3VPN over TE scenario. The networking is shown in Figure 4 below.

CE1 and CE2 belong to the same L3VPN and access the public network through PE1, PE2 and PE3 respectively.

There are two services between CE1 and CE2: voice and OA. The traffic from CE1 to CE2 can be forwarded through two paths: Path1 (PE1->PE2->CE2) and Path2 (PE3->P5->P6->PE2->CE2). Among them, the reliability of path 1 is high and the transmission delay is low. Path 2 has a large bandwidth.

The voice service traffic will be forwarded through Path1 first. The OA service traffic will be forwarded through Path2 first. When the transmission delay of Path1 exceeds the threshold value and Path2 can meet the delay requirements, switch the voice service to Path2.

When the remaining bandwidth of Path2 is less than the bandwidth guarantee threshold, if Path1 still has enough remaining bandwidth, the OA traffic exceeding the bandwidth will be directed to Path1.




Figure 4

The configuration on the head node CE1 includes the following three parts. These configurations can be directly configured on the node or distributed through the controller.

1. Define three Policy Selector policies, and specify the threshold of network quality, path priority and the corresponding path color value for routing.

```
routing-policy-selector irp1
  traffic-delay threshold 1000ms
  priority 1 mapping-to color 100
  priority default mapping-to color 200
routing-policy-selector irp2
  remaining-bandwidth threshold 50M
  priority 1 mapping-to color 200
  priority default mapping-to color 100
```

2. Configure forwarding paths.

```
sr-policy policy-A (color 100, CE2_SID)
  segment-list <SID_PE1, SID_PE2, SID_CE2>
sr-policy policy-B (color 200, CE2_SID)
  segment-list <SID_PE3, SID_P5, SID_P6, SID_PE2, SID_CE2>
```

3. Configure corresponding Policy Selector policies for services with specified characteristics in the parent SRv6 Policy group.

```
parent-sr-policy sr-policy-1(color 10, CE2_SID)
  service voice use routing-policy-selector irp1
  service oa use routing-policy-selector irp2
```

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

This document does not introduce any security considerations.

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/rfc/rfc8402>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/rfc/rfc9256>>.

- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Ed., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/rfc/rfc9830>>.
- [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/rfc/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/rfc/rfc8174>>.

9.2. Informative References

- [I-D.ietf-idr-performance-routing]
Xu, X., Hegde, S., Talaulikar, K., Boucadair, M., Jacquenet, C., and J. Dong, "BGP Performance-aware Routing Mechanism", Work in Progress, Internet-Draft, draft-ietf-idr-performance-routing-05, 5 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-performance-routing-05>>.
- [I-D.cheng-spring-sr-policy-group]
Cheng, W., Wenying, J., Lin, C., Chen, R., Zhang, Y., and Y. Liang, "SR Policy Group", Work in Progress, Internet-Draft, draft-cheng-spring-sr-policy-group-08, 17 June 2025, <<https://datatracker.ietf.org/doc/html/draft-cheng-spring-sr-policy-group-08>>.
- [I-D.liu-ippm-srv6-bandwidth-measurement]
Liu, Y., Lin, C., Qiu, Y., Liu, Y., and Y. Liang, "Measurement Method for Bandwidth of SRv6 Forwarding Path", Work in Progress, Internet-Draft, draft-liu-ippm-srv6-bandwidth-measurement-00, 26 November 2024, <<https://datatracker.ietf.org/doc/html/draft-liu-ippm-srv6-bandwidth-measurement-00>>.

Acknowledgements

The authors would like to thank the following for their valuable contributions of this document.

TBD.

Authors' Addresses

Feng Yang
China Mobile
Beijing
China
Email: yangfeng@chinamobile.com

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