SRops
Internet-Draft
Intended status: Informational
Expires: 13 August 2026

F. Yang
China Mobile
C. Lin
New H3C Technologies

9 February 2026

SR Policy Selector
draft-yang-srv6ops-policy-selector-02

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 SR Policies based on network quality.

Table of Contents

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 SR Policy is associated with one or more candidate paths. A composite candidate path acts as a container for grouping SR Policies. As described in section 2.2 in [RFC9256], the composite candidate path construct enables combination of SR 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 SR Policies. For convenience, the composite candidate path formed by the combination of SR Policies is called parent SR Policy in [I-D.ietf-spring-sr-policy-group].

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

This document describes a policy selection mechanism among the candidate SR 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].

CRM: Customer Relationship Management is a critical application that requires low bandwidth and low latency network connection.

Parent SR Policy: Refer to [I-D.ietf-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 network 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. CRM and conference traffic have different SLA requirement, and expected be carried by different SR Policies. Generally, from CE1 to CE2, conference services with low latency requirements are forwarded along SR Policy PE1->P1->P2->PE2 and PE1->P3->P4->PE2. The CRM traffic is forwarded along the other SR Policy PE1->P5->P6->PE2. When failure or degradation happened in CRM SR Policy, it should be possible to switchover CRM traffic to conference SR Policy.

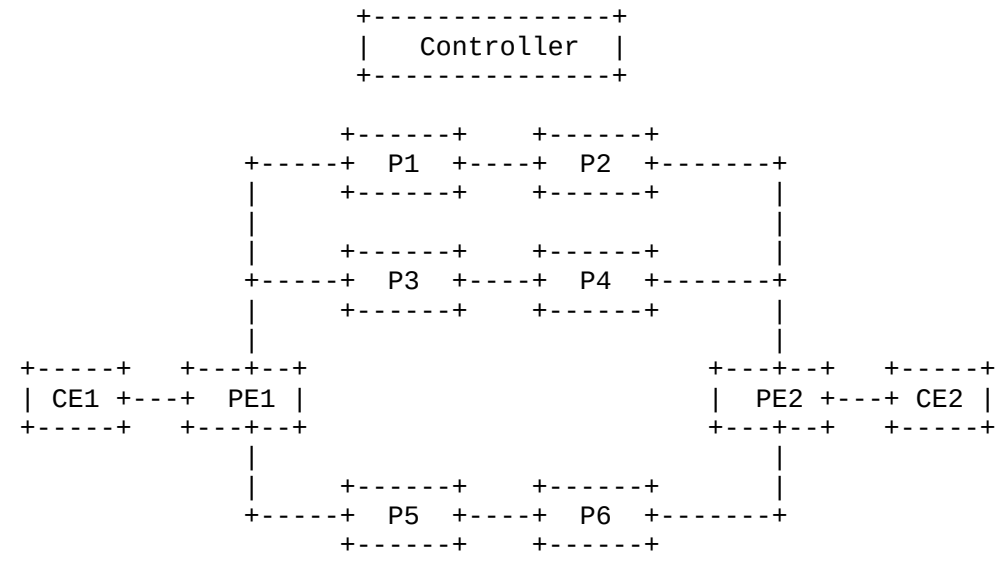

Figure 1

Based on such scenario, the following requirements should be met:

1. Maximize failure/degradation protection

   In case of failure or degradation detected on one SR 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 SR Policy degradation.

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

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

Take Figure 1 as an example, there is a parent SR Policy between PE1 to PE2, which has multiple constituent SR Policies. An SR Policy selection mechanism is needed, which should select best constituent SR Policy in the parent SR Policy. When the head node detects the quality degradation of the active constituent SR Policy, it will select another one in the parent SR Policy.

5. SR Policy Selector

5.1. Processing Model

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

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

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

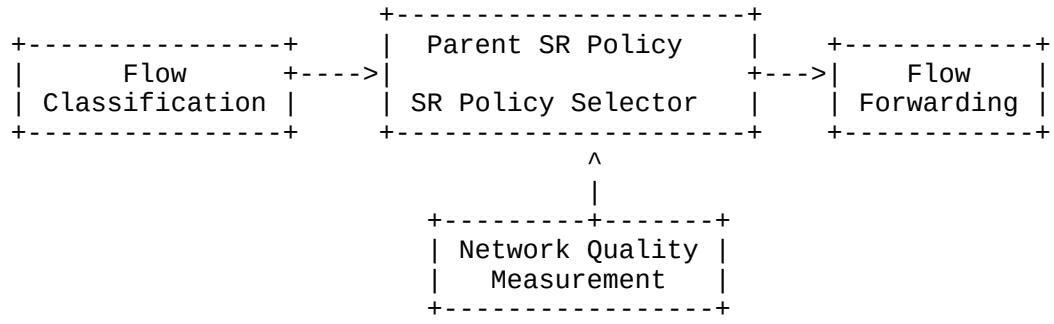

Figure 2

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. SR Policy Selector

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

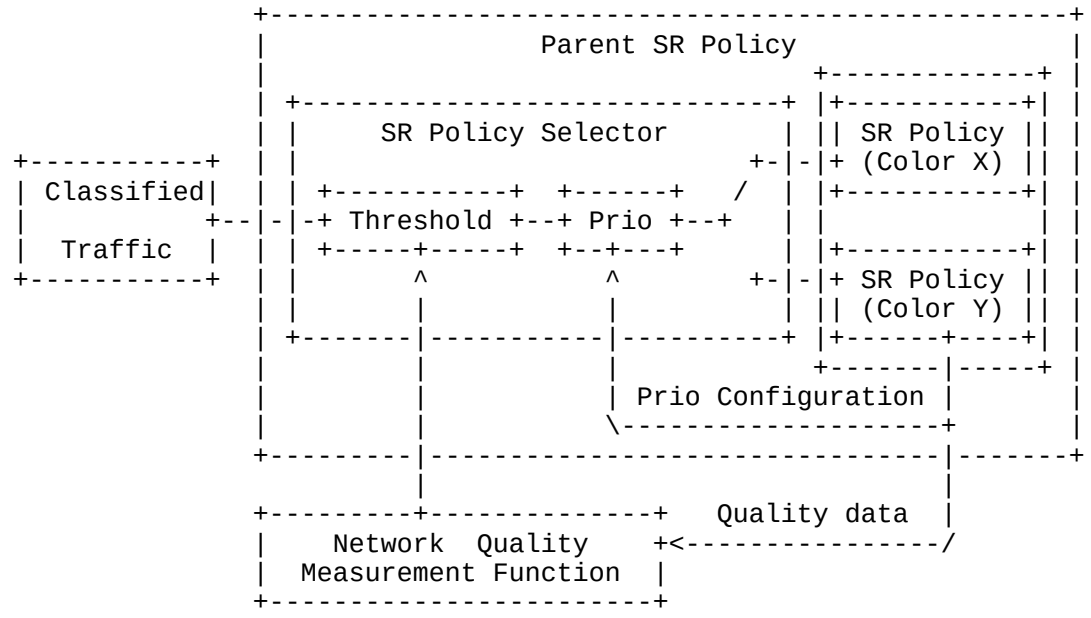

Figure 3

Each parent SR Policy contains multiple constituent SR Policies. Each constituent SR Policy will include two new configuration parameters, "priority" and "threshold" in this proposal. The constituent SR Policy with the highest priority and qualified threshold will be selected to carry the traffic.

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 SR Policy to the one with higher priority.

5.4. 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 SR 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 STAMP, TWAMP, SR 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.

5.5. Flow Forwarding

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

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

6. Examples of SR Policy Selector

The application of SR Policy Selector is described in detail in L3VPN over TE scenario. Take the example shown in Figure 1.

There are two services between CE1 and CE2: conference and CRM. The traffic from CE1 to CE2 can be forwarded through two paths: Path1 (PE1->P1->P2->PE2 and PE1->P3->P4->PE2) and Path2 (PE1->P5->P6->PE2).

The conference service traffic will be forwarded through Path1 first. The CRM 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 conference service to Path2.

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

The configuration on the head node PE1 includes the following three parts. These configurations can be directly configured on the node or distributed through the controller.

1. Configure the parent SR Policy.

```
parent-sr-policy sr-policy-1(color 10, PE2_SID)
  service conference use routing-policy-selector irp1
  service crm use routing-policy-selector irp2
```

2. Configure constituent SR Policy.

```
sr-policy path1 (color 100, PE2_SID)
  segment-list <SID_P1, SID_P2, SID_PE2>
  segment-list <SID_P3, SID_P4, SID_PE2>
sr-policy path2 (color 200, PE2_SID)
  segment-list <SID_P5, SID_P6, SID_PE2>
```

3. Define three SR Policy Selector policies, and specify the threshold of network quality, priority.

```
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
```

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.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>.

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