srv6ops
Internet-Draft
Intended status: Informational
Expires: 1 January 2026

F. Yang
China Mobile
C. Lin
New H3C Technologies

30 June 2025

Intelligent Routing Method of SR Policy
draft-yang-srv6ops-intelligent-routing-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 an intelligent routing method for SR Policy based on network quality in MPLS and IPv6 environments.

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. The ingress node steers packets into a specific path according to the Segment Routing Policy (SR Policy) as defined in [RFC9256]. In order to distribute SR Policies to the headend, [I-D.ietf-idr-segment-routing-te-policy] specifies a mechanism by using BGP.

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.cheng-spring-sr-policy-group].

This document describes an intelligent routing method for SR Policy based on network quality in MPLS and IPv6 environments.

1.1. 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.

2. Terminology

The definitions of the basic terms are identical to those found in Segment Routing Policy Architecture [RFC9256].

3. Problem and Requriements

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 were carried by different SR 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 high bandwidth path PE3->P5->P6->PE2->CE2. When failure or degradation happened in OA traffic SR Policy, there should be possible to assure basic communication for OA traffic by using voice bandwidth.

In single SR 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 SR Policies.

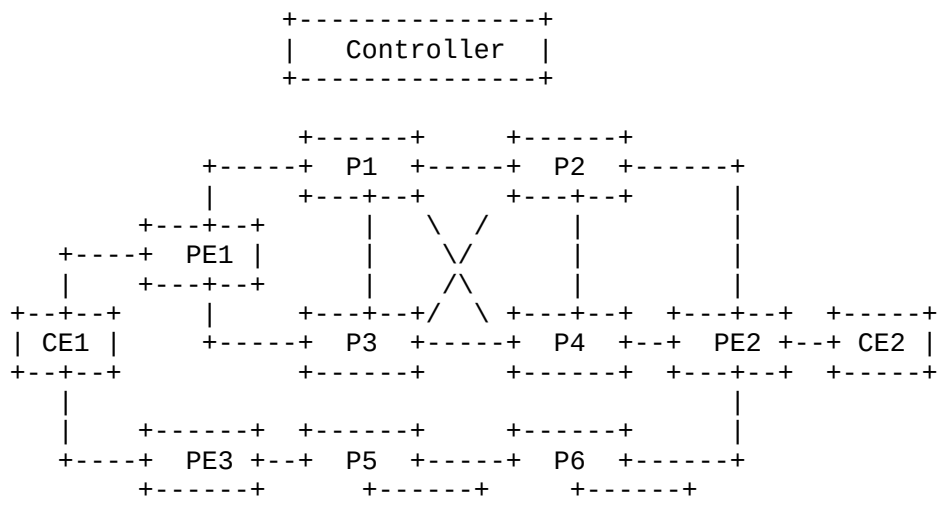

Figure 1

Based on such scenarios, the following requirements are proposed:

1. Maximize failure/degradation protection

   In case of failure or degradation detected on one 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.

In order to better meet these requirements, this document proposes an intelligent routing method for SR policy based on network quality requirement. The head end node selects the optimal path according to the current network quality to improve the path switching speed and forwarding performance.

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 SR Policy group.

Define a parent SR Policy group for the above CE1 to CE2 services. Specify the steering policies of services in the parent SR Policy group. Different services can select different SR Policy paths in the parent SR Policy group according to different quality requirements. When the head node perceives that the quality of the path of a service is deteriorating (such as bandwidth or delay degradation), it searches for other path in the group that is suboptimal and also meets its quality requirements.

4. Intelligent Routing Method for SR Policy

4.1. Processing Model

The path priority is assigned to the SR policy forwarding path manually by the controller. Each path with quality requirement will be assigned with a priority value. The lower the value, the higher the priority. That is, when there is a group of qualified paths, best path will be selected with higher priority.

Configure multiple SR policy paths for the service flows with specified characteristics in the parent SR Policy group. Assign the corresponding path priority to each path according to the priority order of the path. If there is a backup path for the SR policies, lower priority value should be used according to the quality requirements.

After receiving the service packet with the specified characteristics, when the network quality is good, the traffic is forwarded through the path with high priority. When the network quality degradation is happened on the high priority path, such as the packet loss rate exceeds the acceptable range, switch to the next high priority path of the service.

If the quality of the high priority forwarding path is restored and the specified quality requirements are met, the traffic is switched from the low priority forwarding path to the high priority forwarding path after a period of wait-to-restore time.

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

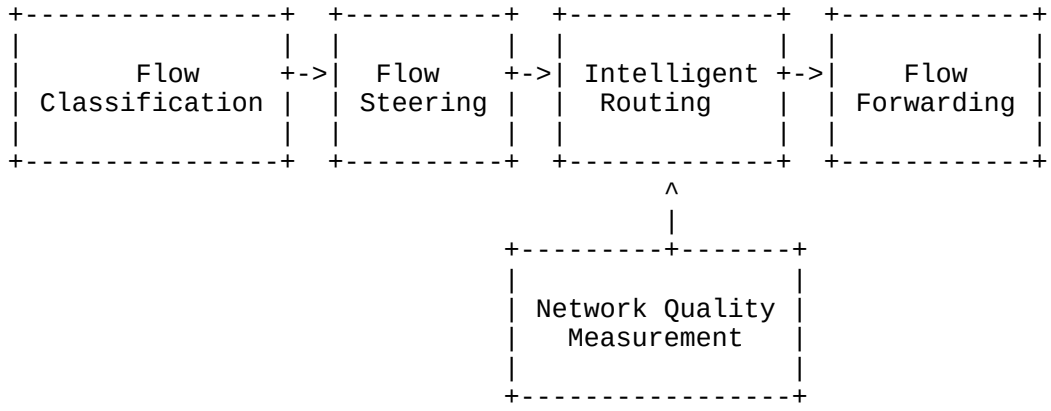

Figure 2

The functions of each unit are described below.

4.2. Flow Classification

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

The head node can match flow characteristics in its ingress interfaces (upon any field such as Ethernet destination/source/VLAN/ TOS or IP destination/source/DSCP or transport ports or application attribute etc.) and color them with an internal per-packet forwarding-class variable.

4.3. Flow Steering

According to the forwarding class variables determined by the Flow classification, the header node selects the matching forwarding path, that is, selects the SR policy or the parent SR policy representing a group of policies.

If multiple SR policy forwarding paths are configured for the traffic flow with the specified characteristics, all valid SR policies will be retrieved and handed over to the Intelligent Routing unit to select the optimal forwarding path.

4.4. Intelligent Routing

According to the SR policy(policies) provided by Flow Steering, the Intelligent Routing unit obtains the current quality of each SR policy path from the Network Quality Measurement unit. Based on the mapping between the quality and the priority of intelligent routing, it selects the forwarding path with the highest priority and the quality measurement of the SR Policy. Only those qulified SR Policies which can reach the threshold are considered as candidate SR Policies.

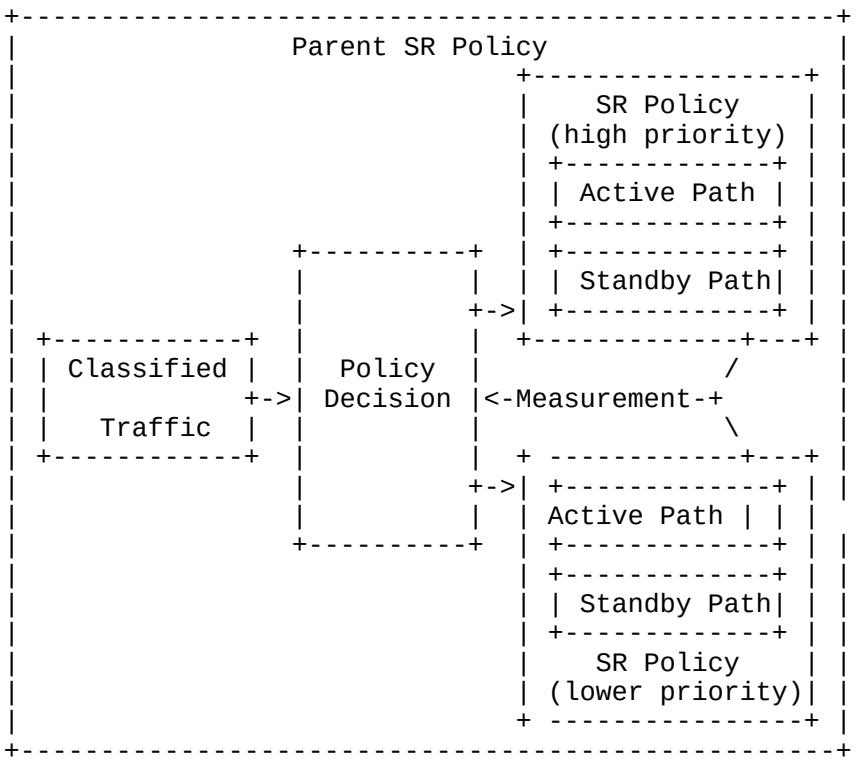

Figure 3

When the network quality is better than the threshold, the traffic is forwarded by the policy with high priority. When the network quality of the high priority SR Policy degrades, such as the loss rate increasion, the Policy Decision module will switch the traffic to the next high priority one in the candidate SR Policies. Similarly, when the next higher priority SR Policy forwarding path cannot meet the forwarding quality requirements, switch to the lower priority SR Policy path.

If the quality of the high priority forwarding path gets better and meets the specified quality requirements, the traffic can be recovered from the low priority forwarding path to the high priority SR Policy forwarding path after a period of wait-to-restore time. The purpose of wait-to-restore time is specified in order to prevent flapping between SR Policies.

To avoid frequent path switching when the network quality is unstable, if the current path can meet the forwarding quality requirements, the head node can choose not to automatically switch back to the higher priority path in case of the quality of the higher priority path is restored. The device can provide a configuration for automatic failback, and add a wait-to-restore timer. Only after automatic restore is allowed and the wait-to-restore timer is timeout, the forwarding path switch from the current path that meets the quality requirements to the path with higher priority.

4.5. 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 Intelligent Routing unit, which decides whether to switch paths.

The following network quality parameters of forwarding path can be used for path scheduling:

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

The quality parameters of network forwarding path can be obtained through active or passive performance measurement methods, such as iOAM, STAMP, TWAMP, 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.

4.6. Flow Forwarding

The service flow is forwarded according to the path determined by the Intelligent Routing 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.

5. Examples of intelligent routing

The application of intelligent routing 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.

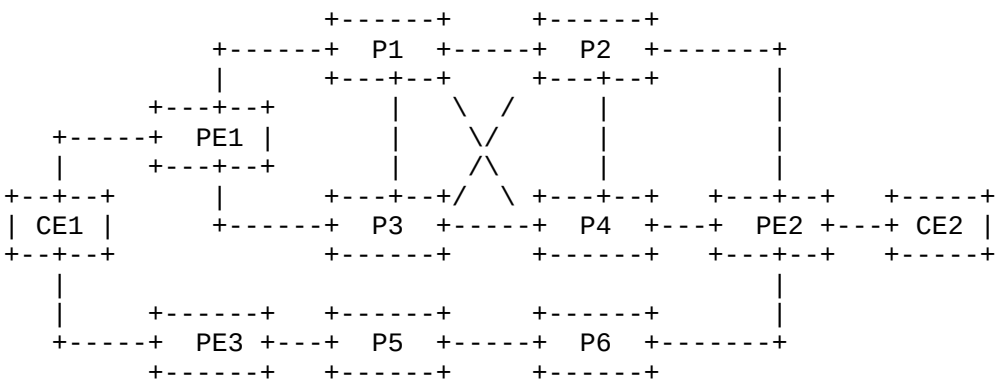

Figure 4

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 intelligent routing policies, and specify the threshold of network quality, path priority and the corresponding path color value for routing.

```
intelligent-routing-policy irp1
  traffic-delay threshold 1000ms
  priority 1 mapping-to color 100
  priority default mapping-to color 200
intelligent-routing-policy 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 intelligent routing policies for services with specified characteristics in the parent SR Policy group.

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

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

This document does not introduce any security considerations.

8. References

8.1. 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>.

[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>.

8.2. Informative References

[I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-26, 23 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-26>.

[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>.

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