

SPRING Working Group
Internet Draft
Intended status: Standards Track
Expires: 12 November 2026

Y. Liu
China Mobile
C. Lin
New H3C Technologies
S. Peng
Huawei Technologies
R. Chen
ZTE Corporation
Z. Ali
Cisco Systems, Inc.
12 May 2026

Flexible Candidate Path Selection of SR Policy
draft-ietf-spring-sr-policy-flexible-cp-selection-00

Abstract

This document describes a flexible method for selecting candidate Segment Routing (SR) policy paths. Based on the real-time resource usage and forwarding quality of candidate paths, the head node can perform dynamic path switching among multiple candidate paths in the SR policy.

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 12 November, 2026.

Copyright Notice

Copyright (c) 2026 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

Liu, et al.

Expires November 2026

[Page 1]

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. Terminology.....	3
3. Background Requirements.....	4
4. Flexible Candidate Path Selection Method.....	5
4.1. Threshold Parameters of Candidate Paths.....	6
4.2. Updated SR Policy Information Model.....	8
4.3. Rules for Setting the Eligibility Attribute.....	9
4.4. Consideration of Multiple Segment Lists.....	10
4.5. Performance Measurement for SR Policy Forwarding Quality...	11
4.6. Flexible Candidate Path Selection Process.....	12
4.7. Delayed Recovery Switch for Candidate Path.....	13
5. Use Cases of Flexible Candidate Path Selection.....	14
5.1. Select the Best Path Based on End-to-End Delay.....	14
5.2. Select the Best Path Based on Available Bandwidth.....	14
5.3. Select the Best Path Based on Actual Bandwidth.....	15
6. IANA Considerations.....	16
7. Security Considerations.....	16
8. References.....	17
8.1. Normative References.....	17
8.2. Informative References.....	17
9. Acknowledgments.....	19
Contributors.....	20
Authors' Addresses.....	21

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].

An SR Policy may have multiple candidate paths that are provisioned or signaled [RFC9830] [RFC8664] from one of more sources. The tie-breaking rules defined in [RFC9256] result in determination of a single "active path" in a formal definition.

According to [RFC9256] the head node must use only the active candidate path for forwarding traffic that is being steered onto a specific policy, except for certain scenarios such as fast reroute where a backup candidate path may be used. A candidate path can be represented as a segment list or a set of segment lists. If a set of segment lists is associated with the active path of a policy, then

the steering of traffic onto the different segment lists is per flow and is based on weighted-ECMP (W-ECMP) according to the relative weight of each segment list.

According to the criteria for the validity of candidate paths described in Section 5 of [RFC9256], if there is a valid segment list in the active candidate path, the active candidate path is valid. When some segment lists of the active candidate path are not valid, the active candidate path may still be valid, but it might not continue to meet the actual forwarding requirements.

[I-D.ietf-spring-sr-policy-eligibility] introduces the concept of an eligibility attribute at the candidate path level, not only at the time of the path computation, but also through topology and network changes to ensure that user intentions are preserved while carrying service traffic.

This document introduces the concept of eligibility refer to [I-D.ietf-spring-sr-policy-eligibility] and specifies a comprehensive quality-driven candidate path switching mechanism. It primarily focuses on defining the precise conditions, operational rules, and dynamic procedures required for performance-driven management. The framework enables autonomous system behavior based on real-time quality assessments, ensuring reliable and adaptive network operations in response to fluctuating performance conditions.

Based on real-time resource usage and the forwarding quality of candidate paths, the head node can dynamically adjust the eligibility attribute value, enabling it to dynamically switch traffic onto different paths among multiple candidate paths within the SR policy.

[RFC2386] provides valuable background on QoS-based routing, details some issues and requirements associated with QoS-based routing, and describes a framework for employing QoS-based routing within the Internet. This document describes an SR Policy mechanism where the traffic is switched between paths based on the resource status of the traversed path. However, it does not address the challenges related to dynamic distributed scheduling or resource reservation along intermediate paths. The document specifies the capability to switch to alternative paths within a strategy when the current path fails to satisfy designated link quality criteria, such as bandwidth, delay, or packet loss. In instances where a controller issues an SR Policy encompassing multiple paths, should a path's link quality not meet the established requirements, a switch to a backup path for forwarding is executed.

2. Terminology

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

3. Background Requirements

When some segment lists of the active candidate path are not valid, according to [RFC9256], if there is a valid segment list in the active candidate path, the active candidate path is still valid. But the paths of remaining segment lists may not meet the SR policy forwarding performance requirements, such as insufficient path bandwidth. Even if there are other candidate paths with lower preference that can meet the forwarding performance requirements in the SR policy, the traffic will continue to be forwarded along the original active candidate path.

As an example, consider the following SR Policy to illustrate the issues present in the current candidate path selection process in detail.

```
SR Policy POL1
Candidate Path CP1
  Preference 200
  Segment List 1 <SID11...SID1i>, Weight 1
  Segment List 2 <SID21...SID2j>, Weight 1
  Segment List 3 <SID31...SID3k>, Weight 1
Candidate Path CP2
  Preference 100
  Segment List 4 <SID41...SID4i>, Weight 1
  Segment List 5 <SID51...SID5j>, Weight 1
  Segment List 6 <SID61...SID6k>, Weight 1
```

There are two static candidate paths CP1 and CP2 in SR policy POL1. CP1 has a higher preference. Both candidate paths are composed of three static segment lists with the same weight. The path indicated by each segment list can carry traffic of 100Mbps bandwidth. When all Segment Lists in CP1 are valid, the effective bandwidth of the candidate path is 300Mbps.

Suppose the bandwidth of the actual traffic forwarded by the SR policy is between 100Mbps and 150Mbps. Because the traffic forwarded on the candidate path will share the load on the three segment list paths according to the weight value, the candidate path can meet the forwarding requirements. The traffic is forwarded on the three segment lists of the higher preference candidate paths of the SR policy.

When segment lists 1 and 2 in the high-preference candidate path CP1 are not valid, according to the candidate path validity criteria described in [RFC9256] Section 5, because segment list 3 in CP1 is still valid, the active candidate path CP1 is still valid. All traffic of SR policy POL1 will continue to be forwarded through the path of CP1. However, because segment list 3 can only forward 100Mbps traffic, over-bandwidth traffic will be discarded.

Of course, when the Segment List path fault is detected, the network device can report the detected fault information to the controller. The controller optimizes the forwarding path after receiving the message. However, this interaction process is relatively long, and it is difficult to meet the requirement for fast switch-over.

When the quality of the high-preference candidate paths deteriorates, due to issues such as insufficient available bandwidth, increased end-to-end transmission delay, or segment lists that fail to meet service requirements, the same need arises. The goal is to switch traffic to other lower preference candidate paths within the SR policy that better satisfy the forwarding quality requirements.

To address this issue, this document proposes a new candidate path selection rule that defines resource thresholds and forwarding quality requirements for candidate paths.

If a candidate path does not satisfy the forwarding quality requirements, its eligibility attribute MUST be set to false. During the active candidate path(CP) selection process, the head-end SHALL use this eligibility attribute as an additional mandatory criterion, in conjunction with the rules defined in [RFC9256], Section 2.9. When a CP's eligibility attribute is false, it indicates that the path cannot forward traffic meeting the specified quality requirements and therefore MUST NOT be considered for active CP selection.

4. Flexible Candidate Path Selection Method

As described in [RFC9256], the candidate path selection process operates primarily on the candidate path Preference. A candidate path is selected when it is valid and it has the highest Preference value among all the valid candidate paths of the SR Policy.

[I-D.karboubi-spring-sr-policy-eligibility] introduces a new attribute at the candidate path level called Eligibility. Only candidate paths with Eligibility as true are considered as part of the active candidate path selection defined in [RFC9256].

This document describes using forwarding quality requirements and resource requirements of candidate paths as eligibility criteria for path selection.

A headend may be informed about the forwarding quality requirements of a candidate path for an SR Policy <Color, Endpoint> through various means, including configuration, PCEP, or BGP. The extensions of BGP and PCEP are described in [I-D.liu-idr-bgp-sr-policy-cp-threshold] and [I-D.liu-pce-sr-policy-cp-threshold].

When a candidate path fails to meet forwarding quality requirements, its Eligibility attribute SHOULD be set to false, thereby excluding it from active candidate path selection.

For candidate paths containing multiple segment lists:

- If a segment list fails to meet forwarding quality requirements, it SHOULD be excluded from forwarding operations.
- When all segment lists under a candidate path fail to meet forwarding quality requirements, the path's Eligibility attribute SHOULD be set to false, disqualifying it from active candidate path selection.

4.1. Threshold Parameters of Candidate Paths

The threshold parameters of candidate paths can include, but are not limited to, the following:

- * Jitter
- * Latency
- * Packet loss

When the jitter, delay, or packet loss of a valid segment list does not meet the specified threshold requirements, the segment list will be deemed not valid and will no longer participate in load sharing traffic.

- * Available bandwidth

$$\text{CP available bandwidth} = \text{CP preset bandwidth} * (\text{Sum of weights of Segment Lists in Up state} / \text{Sum of all Segment List weights})$$

- * Actual bandwidth

The actual bandwidth refers to the sum of the actual available remaining bandwidth of each valid segment list in the candidate path.

Due to the different congestion conditions of each node on the forwarding path, the actual bandwidth that can forward service packets may differ from the preset bandwidth. By utilizing some measurement mechanisms, the actual minimum available bandwidth and actual minimum remaining bandwidth of all nodes along the

path can be obtained. The specific measurement mechanism is not within the scope of this document.

- * Precision Availability Metrics (PAM)

Consider a candidate path of SR policy as a Service Level Objective (SLO) [RFC9543], based on the Precision Availability Metrics (PAM) defined in [RFC9544], determine whether the candidate path meets the forwarding requirements.

If one or more segment lists in the candidate path fail to meet the thresholds, this indicates that these segment lists cannot provide forwarding capabilities that meet the Service Level Agreement (SLA) requirements. These segment lists will be marked as unavailable and will no longer participate in packet forwarding. After excluding these segment lists, the candidate path should be re-verified to determine whether it still meets the forwarding quality requirements. If it does, traffic can continue to be forwarded along that candidate path.

For example, two threshold parameters, delay and available bandwidth, are specified for the candidate path with multiple segment lists. When the delay of a segment list exceeds the threshold, the following processing is performed:

- 1) Remove the segment list from the forwarding path.
- 2) Calculate the current available bandwidth of the CP based on the weight ratio of the remaining effective segment lists and the bandwidth of the CP.
- 3) Check whether the current available bandwidth of the CP still meets the bandwidth threshold requirements.
 - * If the available bandwidth still meets the requirements, the candidate path still meets the forwarding quality requirements, and the traffic is still forwarded along this candidate path.
 - * Otherwise, set the Eligibility attribute of this CP to false. The system will then consider switching service traffic to another active candidate path with better forwarding quality.

If the candidate path does not specify any threshold parameters, select the primary candidate path according to the selection method defined in [RFC9256].

By default, there is no threshold parameter specified on the candidate path.

4.2. Updated SR Policy Information Model

This document defines a quality-driven information model for Segment Routing (SR) Policy. Specifically, it introduces threshold-based performance parameters for forwarding quality under each Candidate Path. When a Candidate Path fails to meet the specified quality thresholds, its Eligibility attribute SHOULD be set to false, thereby excluding it from being selected as the active path.

In summary, the information model enables dynamic and performance-aware SR Policy enforcement, ensuring that only qualified paths are used based on real-time forwarding quality.

SR Policy POL1: <Headend = H1, Color = 1, Endpoint = E1>

 Candidate Path CP1 <Protocol-Origin, Originator,Discriminator>

 Binding SID

 Preference

 Priority

 Threshold parameters of forwarding quality

 Jitter 1

 Latency 1

 Available Bandwidth 1

 Actual Bandwidth 1

 Segment List 1

 Weight W1

 Segments <SID11...SID1i>

 Segment List 2

 Weight W2

 Segments <SID21...SID2i>

 Candidate Path CP2 <Protocol-Origin, Originator,Discriminator>

 Binding SID

 Preference

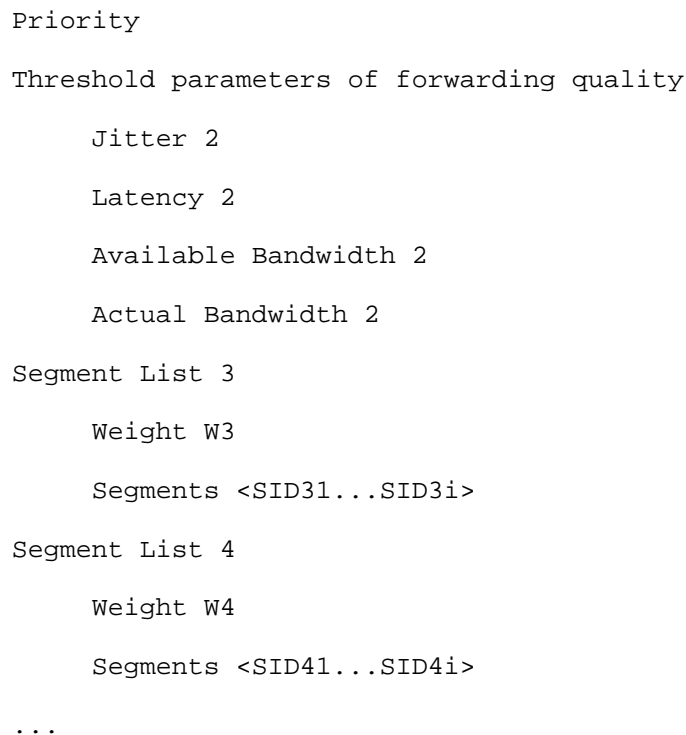


Figure 1: Information Model of SR Policy with forwarding quality

4.3. Rules for Setting the Eligibility Attribute

When a candidate path's current forwarding quality meets the specified threshold requirements, its Eligibility attribute MUST be set to true, indicating this path is valid for:

- * Traffic forwarding operations.
- * Active candidate path selection (per [RFC9256] selection methodology)

Conversely, when a candidate path fails to meet quality requirements, its eligibility attribute MUST be set to false.

For candidate paths containing multiple segment lists:

- If a segment list fails to meet forwarding quality requirements, it SHOULD be excluded from forwarding operations.
- When all segment lists under a candidate path fail to meet forwarding quality requirements, the path's Eligibility attribute

SHOULD be set to false, disqualifying it from active candidate path selection.

For candidate paths without defined threshold parameters:

- * The eligibility attribute MUST default to true.
- * Primary path selection follows [RFC9256] procedures.

When multiple eligible candidate paths coexist in an SR policy:

- * Paths with Eligibility=false MUST NOT participate in active path selection.
- * Detailed behavior is specified in [I-D.ietf-spring-sr-policy-eligibility].

4.4. Consideration of Multiple Segment Lists

When one or more segment lists in the candidate path fail to meet the threshold parameters of the candidate path, it indicates that these segment lists cannot provide forwarding capabilities that meet the SLA requirements. These segment lists will be marked as unavailable and will no longer participate in packet forwarding. After excluding these segment lists, it should be verified whether the candidate path still meets the forwarding quality requirements. If it does, traffic can continue to be forwarded along that candidate path.

For example, two threshold parameters, delay and available bandwidth, are specified simultaneously for the candidate path with multiple segment lists. When the delay of a segment list exceeds the threshold, the following processing is performed:

- 1) Remove the segment list from the forwarding path first.
- 2) Calculate the current available bandwidth of CP based on the weight ratio of the remaining effective segment lists and the bandwidth of CP.
- 3) Check whether the current available bandwidth of CP still meets the bandwidth threshold requirements.
 - * If the available bandwidth still meets the requirements, the candidate path still meets the forwarding quality requirements, and the traffic is still forwarded along this candidate path.
 - * Otherwise, the path's Eligibility attribute SHOULD be set to false, disqualifying it from active candidate path selection.

4.5. Performance Measurement for SR Policy Forwarding Quality

A comprehensive SR Performance Measurement toolset is an essential requirement for measuring network performance to provide Service Level Agreements (SLAs). The following lists several measurement methods for reference; detailed measurement methods are beyond the scope of this document.

- * jitter, delay, or packet loss

To measure jitter, delay, or packet loss of an SR policy candidate path, the Simple Two-Way Active Measurement Protocol (STAMP) can be employed. [I-D.ietf-spring-stamp-srpm-mpis] and [I-D.ietf-spring-stamp-srpm-srv6] describe the performance measurement procedures in SR networks using STAMP as defined in [RFC8762], along with its optional extensions from [RFC8972] and [RFC9503]. Notably, [RFC9503] defines an extended TLV capability which can carry SR path information, ensuring that both the forward and return probe packets follow the same SR path for measurement. To guarantee path symmetry, a return-path SID identifying the path segment can be included within Sub-TLVs. The reflector then uses this return-path SID to construct the SID List for the return probe packet. Specifically, [RFC9503] specifies that the Return Path SRv6 Segment List Sub-TLV may include an identifier for the return path Segment List of an SRv6 Path Segment.

- * Available bandwidth

The Available Bandwidth of a Candidate Path refers to the bandwidth allocated to the Candidate Path during SR Policy path computation based on user requirements. The Segment List weight ratio represents the load sharing proportion of each Segment List under the Candidate Path. The product of the Candidate Path bandwidth and the effective Segment List weight indicates the Available Bandwidth that the Candidate Path can currently carry relative to the user's bandwidth requirements.

Candidate Path Available Bandwidth = Candidate Path Bandwidth * (Sum of UP Segment List Weights / Total Segment List Weights).

If the Available Bandwidth of the currently active Candidate Path cannot meet the requirements, the SR Policy can switchover to a backup Candidate Path that satisfies the bandwidth requirements.

- * Actual bandwidth

By utilizing some measurement mechanisms, the actual minimum available bandwidth and actual minimum remaining bandwidth of all nodes along the SR Policy path can be obtained. One possible implementation to obtain the actual minimum available and remaining bandwidth along a path is as follows:

- 1) The head node sends STAMP probe packets with a DOH header added outside the SRH for SRv6 Policy to record the path's minimum bandwidth.
- 2) At each hop, the egress interface's real-time bandwidth is compared with the minimum value recorded in the DOH header.
- 3) If the interface bandwidth is greater, the DOH field remains unchanged; if it is smaller, the DOH field is updated accordingly.
- 4) Finally, the tail node records the final minimum bandwidth from the DOH header and returns this information to the head node via an extended TLV in the STAMP reply packet.

In operational deployments, real-time bandwidth may exhibit continuous fluctuations. To mitigate CP switchovers resulting from such variability, it is advisable to define a measurement interval for bandwidth monitoring. During this interval, multiple samples SHOULD be collected and averaged to smooth transient anomalies and prevent spurious CP transitions.

It should be noted that the specific measurement mechanism itself is outside the scope of this document.

4.6. Flexible Candidate Path Selection Process

The process of selecting the best candidate path for SR policy through the threshold parameter of the candidate path is as follows.

- 1) Configure the threshold parameters on the candidate path of the head node through manual configuration or controller distribution.
- 2) The head node monitors whether the available resources and forwarding quality of the SR policy candidate path exceed the thresholds. The forwarding quality of path can be obtained through active or passive performance measurement methods, such as iOAM [RFC9378], STAMP [I-D.ietf-spring-stamp-srpm-mpls] and [I-D.ietf-spring-stamp-srpm-srv6], TWAMP [RFC5375], etc. The real-time quality data can be calculated by the controller and distributed to the head node, or calculated by the head node according to the network measurement data. The measurement method and quality data acquisition method are beyond the scope of this document.
- 3) According to the rules described in Section 4.3, if the available resources of the active candidate path fall below the threshold, or if its forwarding quality fails to meet the required threshold, the head node shall select a new active candidate path. For a candidate path comprising multiple segment lists, if

all of its segment lists fail to satisfy the forwarding quality requirements, the path's Eligibility attribute MUST be set to false, thereby excluding it from selection as the active candidate path.

- 4) After the fault on the old active candidate path is repaired or the forwarding quality improves, whether to revert to the previous active candidate path can be specified by the configuration. If fault recovery is required, start a wait timer for delayed recovery. When the timer expires and if the old active candidate path still meets the threshold requirements, the traffic will be switched back to the old higher preference candidate path.

4.7. Delayed Recovery Switch for Candidate Path

To avoid frequent path switching (flapping), both over-threshold switching and fault recovery should be delayed. The delay interval can be adjusted through configuration.

When the quality of service (QoS) of a candidate-path degrades and fails to meet the threshold requirement, the candidate-path will be marked as invalid and will not carry traffic.

When the QoS of a candidate-path recovers and meets the threshold requirement, it will wait for the Wait-to-Restore (WTR) period. If the QoS does not degrade during this time, the candidate-path will resume carrying traffic. Otherwise, it will be marked as invalid again.

Taking the WTR of 30 seconds as an example, the explanation is as follows:

0(x)	10(WTR)	35(x)	40(WTR)	70(Recovery)	90(x)	95(WTR)	125(Recovery)
------	---------	-------	---------	--------------	-------	---------	---------------

--	--	--	--	--	--	--	--

At 0 seconds, the candidate-path degraded and was set to invalid.

At 10 seconds, the candidate-path recovered, and the WTR timer was started, waiting for 30 seconds.

At 35 seconds, the candidate-path degraded again, and the waiting timer was canceled.

At 40 seconds, the candidate-path recovered, and the WTR waiting timer was started, waiting for 30 seconds.

At 70 seconds, during the wait-to-restore period, the candidate-path did not degrade; thus, the WTR wait ended and it was restored to valid.

At 90 seconds, the candidate-path degraded again and was set to invalid.

At 95 seconds, the candidate-path recovered, and the WTR waiting timer was started, waiting for 30 seconds.

At 125 seconds, during the wait-to-restore period, the candidate-path did not degrade; thus, the WTR wait ended and it was restored to valid.

5. Use Cases of Flexible Candidate Path Selection

The example SR policy described in Section 3 is used in the sections that follow to illustrate how the flexible candidate path selection method switches between candidate paths.

SR policy POL1 has two candidate paths CP1 and CP2. The Preference of CP1 is 200, and the Preference of CP2 is 100. Both candidate paths are composed of three segment lists with the same weight.

5.1. Select the Best Path Based on End-to-End Delay

The quality requirement for the services carried on the SR policy is that the transmission delay must be less than 200ms. The bandwidth of the actual traffic forwarded by the SR policy is between 100Mbps and 150Mbps.

When the delay of Segment List 1 does not meet the requirements, the head node continues to check the available bandwidth of CP1. Due to segment list 2 only having 100Mbps bandwidth, it cannot meet the actual traffic forwarding requirements. CP1's Eligibility attribute is set to false, triggering the selection of CP2 as POL1's new active candidate path. The traffic forwarded by POL1 is switched to the path of CP2 for forwarding.

```
SR Policy POL1
Candidate Path CP1
  Preference 200
  Delay threshold 200ms //Delay<=200ms
  Segment List 1 <SID11...SID1i>, Weight 1 //100M, Delay>1s
  Segment List 2 <SID21...SID2i>, Weight 1 //100M, Delay<100ms
Candidate Path CP2
  Preference 100
  Delay threshold 200ms //Delay<=200ms
  Segment List 3 <SID31...SID3i>, Weight 1 //100M, Delay<100ms
  Segment List 4 <SID41...SID4i>, Weight 1 //100M, Delay<100ms
```

5.2. Select the Best Path Based on Available Bandwidth

The preset bandwidth for both CP1 and CP2 is 300Mbps. Each segment list can carry a maximum of 100Mbps traffic. The quality requirement for service traffic is that the available bandwidth of the forwarding path must not be less than 150Mbps.

```
SR Policy POL1
Candidate Path CP1
  Preference 200
  Preset bandwidth 300Mbps
  Available bandwidth threshold 150Mbps
  Segment List 1 <SID11...SID1i>, Weight 1
  Segment List 2 <SID21...SID2j>, Weight 1
  Segment List 3 <SID31...SID3k>, Weight 1
Candidate Path CP2
  Preference 100
  Preset bandwidth 300Mbps
  Available bandwidth threshold 150Mbps
  Segment List 4 <SID41...SID4i>, Weight 1
  Segment List 5 <SID51...SID5j>, Weight 1
  Segment List 6 <SID61...SID6k>, Weight 1
```

First, take the available bandwidth as the threshold parameter of POL1. The threshold for configuring the available bandwidth is 150Mbps. When the available bandwidth of the candidate path is less than 150Mbps, perform path switching.

Normally, the three segment lists of CP1 and CP2 are valid. The available bandwidth of CP1 is 300Mbps, and the available bandwidth can meet the threshold requirements. CP1's Eligibility attribute is set to true, CP1 is selected as the active candidate path according to the Preference.

If the paths indicated by Segment List 1 and 2 fail, Segment List 1 and 2 become not valid, and the available bandwidth of CP1 becomes 100Mbps. Because the available bandwidth of CP1 is lower than the specified threshold, CP1 has failed to meet the forwarding quality requirements. CP1's Eligibility attribute is set to false. There is a need to reselect the active candidate path for POL1.

The three segment lists of the low-preference candidate path CP2 of POL1 are valid, and the available bandwidth can meet the threshold requirements. CP2's Eligibility attribute is set to true. CP2 is selected as the new active candidate path of POL1. The traffic forwarded by POL1 will be switched to the path of CP2 for forwarding.

5.3. Select the Best Path Based on Actual Bandwidth

In scenarios involving the actual available bandwidth measurement method for SRv6, as described in

[I-D.liu-ippm-srv6-bandwidth-measurement], the quality requirement for the services carried on the SR policy mandates that the actual available bandwidth of the forwarding path must exceed 80 Mbps. If traffic congestion occurs on a node in Segment List 1, resulting in a maximum forwarding capacity of only 50 Mbps for service traffic,

and if Segment List 2 is either in a down state or has exceeded the delay threshold, Segment List 2 will not participate in load sharing traffic.

When the aggregate available bandwidth of CP1 falls below 80 Mbps:

- * CP1's eligibility attribute is set to false.
- * CP2's eligibility attribute is set to true (provided it meets forwarding requirements).
- * CP2 SHALL become POL1's new active candidate path.

SR Policy POL1

Candidate Path CP1

Preference 200

Preset bandwidth 200Mbps

Actual available bandwidth threshold 80Mbps

Segment List 1 <SID1l...SIDli>, Weight 1

(Actual available bandwidth is only 50Mbps.)

Segment List 2 <SID2l...SID2j>, Weight 1

(In Down state, or the delay has exceeded the threshold.)

Candidate Path CP2

Preference 100

Preset bandwidth 300Mbps

Actual available bandwidth threshold 80Mbps

Segment List 3 <SID4l...SID4i>, Weight 1 (100Mbps)

Segment List 4 <SID5l...SID5j>, Weight 1 (100Mbps)

Segment List 5 <SID6l...SID6k>, Weight 1 (100Mbps)

The traffic forwarded by POL1 will switch to the path of CP2 for forwarding.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

[RFC8754] defines the notion of an SR domain and use of SRH within the SR domain. Procedures for securing an SR domain are defined the section 5.1 and section 7 of [RFC8754]. This document does not impose any additional security challenges to be considered beyond security threats described in [RFC8754], [RFC8679] and [RFC8986].

The traffic switchover mechanism defined in this document, such as the ability to forcibly switch traffic from one control plane to another, may redirect traffic to an attacker's preset path. Additionally, switching traffic to another CP could overload network resources, leading to service unavailability or operational

failures. Similarly, frequent flapping during switchovers may compromise network stability. Therefore, it is essential to ensure that this SR network operates within a trusted security domain while implementing safeguards like proper configuration and delayed switchback mechanisms to maintain secure SR Policy operation.

8. References

8.1. Normative References

- [I-D.ietf-spring-sr-policy-eligibility] Karboubi, A., Shah, H., Stone, A., Schmutzer, C. and Maheshwari, P., "Eligibility Concept in Segment Routing Policies", draft-ietf-spring-sr-policy-eligibility-00 (work in progress), January 2026.
- [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>>.
- [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/info/rfc8402>>.
- [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>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., Hardwick, J., "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/info/rfc9830>>.

8.2. Informative References

- [I-D.liu-idr-bgp-sr-policy-cp-threshold] Liu, Y., Lin, C., Qiu, Y., " BGP Extension for Distributing CP Threshold Constraints of SR Policy", draft-liu-idr-bgp-sr-policy-cp-threshold-02 (expired), November 2024.

- [I-D.liu-pce-sr-policy-cp-threshold] Liu, Y., Lin, C., Qiu, Y., "PCEP Extension to Support Signaling Candidate Path Threshold Constraints of SR Policy", draft-liu-pce-sr-policy-cp-threshold-03 (expired), February 2025.
- [I-D.liu-ippm-srv6-bandwidth-measurement] Liu, Y., Lin, C., Qiu, Y., Liu, Y., Liang, Y., "Measurement Method for Bandwidth of SRv6 Forwarding Path", draft-liu-ippm-srv6-bandwidth-measurement (expired), November 2024.
- [I-D.ietf-spring-stamp-srpm-mpls] Gandhi, R., Filsfils, C., Janssens, B., Chen, M., and Foote, R., "Performance Measurement Using Simple Two-Way Active Measurement Protocol (STAMP) for Segment Routing over the MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-mpls-01, 2 April 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-mpls-01>>.
- [I-D.ietf-spring-stamp-srpm-srv6] Gandhi, R., Filsfils, C., Janssens, B., Chen, M., and Foote, R., "Performance Measurement Using Simple Two-Way Active Measurement Protocol (STAMP) for Segment Routing over the IPv6(SRv6) Data Plane", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-srv6-01, 2 April 2026, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-srv6-01>>.
- [RFC2386] Crawley, E., Nair, R., Rajagopalan, B. and H. Sandick, "A Framework for QoS-based Routing in the Internet", RFC 2386, August 1998.
- [RFC5375] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., Babiarz, J., "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5375, DOI 10.17487/RFC5375, October 2008, <<https://www.rfc-editor.org/info/rfc5375>>.
- [RFC9378] Brockners, F., Bhandari, S., Bernier, D., Mizrahi, T., "In Situ Operations, Administration, and Maintenance (IOAM) Deployment", RFC 9378, DOI 10.17487/RFC9378, April 2023, <<https://www.rfc-editor.org/info/rfc9378>>.
- [RFC9543] Farrel, A., Ed., Drake, J., Ed., Rokui, R., Homma, S., Makhijani, K., Contreras, L., and J. Tantsura, "A Framework for Network Slices in Networks Built from IETF Technologies", RFC 9543, DOI 10.17487/RFC9543, March 2024, <<https://www.rfc-editor.org/rfc/rfc9543>>.

[RFC9544] Mirsky, G., Halpern, J., Min, X., Clemm, A., Strassner, J., Francois, J., "Precision Availability Metrics for Services Governed by Service Level Objectives (SLOs)", RFC 9544, DOI 10.17487/RFC9544, March 2024, <<https://www.rfc-editor.org/info/rfc9544>>.

9. Acknowledgments

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

TBD

Contributors

Gyan S. Mishra
Verizon Inc.
Email: gyan.s.mishra@verizon.com

Yuanxiang Qiu
New H3C Technologies
Beijing
China
Email: qiuyuanxiang@h3c.com

Authors' Addresses

Yisong Liu
China Mobile
Beijing
China
Email: liuyisong@chinamobile.com

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

Shuping Peng
Huawei Technologies
Beijing
China
Email: pengshuping@huawei.com

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

