

SPRING Working Group
Internet-Draft
Updates: 9256 (if approved)
Intended status: Standards Track
Expires: 13 July 2026

A. Karboubi, Ed.
H. Shah, Ed.
Ciena
A. Stone
Nokia
C. Schmutzer
Cisco
P. Maheshwari
Airtel India
9 January 2026

Eligibility Concept in Segment Routing Policies
draft-ietf-spring-sr-policy-eligibility-00

Abstract

Segment Routing (SR) introduces new challenges for pinning candidate paths on their intended paths (the path the PCE computed based on provided intent and may have made bandwidth reservations on). The actual path through a network can change or no longer meet the required constraints if a SID list of an SR Policy candidate path is not fully expressed as a list of adjacency SIDs or when a change in the topology does happen. The introduction of the new candidate path eligibility concept permits a path to be signaled and established as operationally up, but controls whether the path is eligible to carry traffic, thus influencing its active state. The eligibility concept allows a system (operator, pce, headend, etc.) to set eligibility as false when path deviations may have occurred, or path constraints are no longer met for one or more SID lists of a candidate path and clear it when candidate path deviations are removed or constraints are met again.

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 13 July 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 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
1.1. Requirements Language	3
2. Terminology	3
3. Problem statement and illustrative examples:	4
3.1. Example 1 : Deviation from intent due to failures:	4
3.2. Example 2 : Delay sensitive paths:	6
4. The eligibility concept	7
5. Protocol and model changes	7
5.1. Active candidate path selection algorithm	8
5.2. PCEP extensions	8
5.3. SR policy Yang changes	8
5.4. BGP	8
6. IANA considerations	8
7. Security considerations	8
8. References	8
8.1. Normative References	9
8.2. Informative References	9
Acknowledgements	10
Contributors	10
Authors' Addresses	10

1. Introduction

Service providers require that services are delivered on traffic engineered transport such as SR enabled network. This requires path computations carried out by PCE or ingress PE based on operator defined constraints that reflect the service level agreements (SLAs) provided to client. The examples of such constraints are guaranteed bandwidth, end-to-end delay or other topology constraints.

This document introduces a concept of an eligibility attribute at the candidate path level, not only at the time of the computation but also through topology and network changes to ensure that user intentions are preserved while carrying service traffic. The eligibility attribute of the candidate path is then used as an additional mandatory criteria by the head-end during the selection process of active CP in addition to rules specified in [RFC9256] section 2.9. For example, there could exist a candidate path with highest preference, with validated SID list that is operationally up, and OAM monitored but not eligible for selection as active path based on eligibility attribute set to false.

Note that this document focuses on introduction of eligibility concept, and not necessarily the in-depth use cases description, the criteria that should alter the eligibility of a candidate path nor reasons why a system may want to keep a path operationally up yet prevent it from carrying client traffic; these shall be described in their appropriate use case document to detail the reasons and behavior for setting and clearing the path eligibility. It also worth noting that eligibility of a path may be set/unset by various actors and various conditions. (e.g. ingress PE setting path as ineligible based on S-BFD and PCE setting it as eligible based on link recovery or other condition). We present some examples and use cases in Section 3.

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

SID : Segment Identifier

SLA : Service Level Agreement

SR : Segment Routing

CS-SR : Circuit-Style Segment Routing

PCE : Path Computation Element

PCEP : Path Computation Element Communication Protocol

3. Problem statement and illustrative examples:

The general purpose of eligibility is to keep an SR Policy Candidate Path operationally signaled and operationally up including any control plane, assurance, or measurement-related functions while preventing active service or upstream traffic from traversing it. This allows the path to be excluded from the user data plane, when necessary, while still maintaining its signaling and associated measurements. Without keeping the Candidate Path Up it is possible various measurements or operational status cannot be monitored and evaluated. The operationally up nature permits measurements to still continue which can help determine when the path is ready to resume forwarding user traffic and eligibility to be re-enabled.

The following sections describe some example scenarios that have value with the eligibility concept.


3.1. Example 1 : Deviation from intent due to failures:

A PCE computes a path for the service according to the network state and available capacity at that time. These paths are referred to as intended paths. It then encodes the intended path into SIDs using a combination of node and adjacency SIDs. Nodes in the network forward packet to node SID N by using their IGP (or flex-algo) shortest paths to N. This is referred to as path expansion. At the time of installing the SID list, this expansion and the intended path are identical.

However, network changes, particularly link and/or node failures may cause the intended path and this path expansion to deviate resulting in a service traffic to use resources on a path that the PCE did not reserve any bandwidth on, causing service degradation for both this service and the other services on that path. Note that BW is given here as a constraint example only, the deviation could be causing longer delays or violating other service based constraints.

Both the failure and repair cases are illustrated using the example network topology of figure 1. An SR Policy from node A to node Z with two diverse traffic engineered candidate paths was computed by PCE and signaled to head end node A resulting in the following intended paths and their respective SID List:

- * Candidate path 1: intended path A-B, B-D, D-E, E-Z links and signaled as SID list B, E, Z
- * Candidate path 2: intended path A-C, C-D, D-F, F-Z links and signaled as SID list C, F, Z




```

SR Policy A-Z:
  Candidate path1
    SIDList1 [B,E,Z]
  Candidate path2
    SIDList2 [C,F,Z]

```

Figure 1: SR policy with 2 diverse candidate paths

In Figure 2, link B-D fails. The expected behavior is to start using the second candidate path. Though this path may be used initially, once the IGP converges, the candidate path 1 becomes valid as node B regains a shortest path to the next node SID E. Once the headend switches to the candidate path 1, the intended path and the expansion of the SID list which now becomes (A-B, B-G, G-D, D-E, E-Z) deviate. The service starts to use resources on B-G and G-D links where the PCE has not made a bandwidth reservation.



SR Policy A-Z:

Candidate path1

SIDList1 [B,E,Z] --> deviation from intended path due to failure

Candidate path2

SIDList2 [C,F,Z]

Figure 2: SR policy CP1 deviation after link failure and IGP convergence

This document proposes a simple extension to the active candidate path selection algorithm defined in [RFC9256] which renders the candidate path 1 ineligible for selection at the head-end node when system determines that traffic shall not be using this path even if it seems valid.

In the example above, a system could set the CP1 eligibility as false when it detects path failure via some CCV mechanism (e.g. S-BFD, STAMP, etc.) rendering path ineligible for selection, the path may become operationally up after IGP convergence, but it will remain unavailable for selection until the eligibility is cleared.

3.2. Example 2 : Delay sensitive paths:

Using same policy example illustrated in figure 1, the policy could have a constraint to not use a path when its end-end delay exceeds a given value D1. A link B-D for example while still up, could have its delay value increased so that overall policy delay now exceeds D1. The expected behavior is to start using the second candidate path as its delay is meeting the original constraint. In this case, a system could set the eligibility as false when it detects that path delay exceeds D1 (e.g. using STAMP) rendering path ineligible for selection, and because the path is still operationally up and monitored by STAMP, when the delay condition clears, the system could

clear the eligibility for the monitored path.

Note that the above examples are for illustration purposes only, The entities acting on the eligibility and its conditions are outside the scope of this document and would be covered under separate use case documents such as [I-D.karboubi-spring-sidlist-optimize-cs-sr]. Note that it is important to keep the path operationally up and under the purview of any OAM/CCV as we may rely on OAM protocol (e.g. STAMP measuring e2e delay) to determine the eligibility of the CP.

4. The eligibility concept

We introduce a new attribute at the candidate path level called eligibility. Candidate path selection logic is modified so that eligibility must be considered as part of the active candidate path selection defined in [RFC9256]; that is, only candidate paths with eligibility as true, must be considered for carrying traffic.

The eligibility of a path can be controlled by head end, a PCE or user, this is outside the scope of this document, but one such use case is defined under [I-D.karboubi-spring-sidlist-optimize-cs-sr].

Usually marking a path as ineligible can be triggered by a distinct set of conditions (e.g. delay OR path deviation) and the responsibility for setting ineligibility can be split amongst different components, but it is advisable that the clearing of eligibility is ideally performed by a single component having visibility of all conditions (user intent) and not split it amongst distinct components as all conditions need to be met prior to marking path as eligible again.

In case an implementation or use case requires the clearing of eligibility to be also split between distinct components care needs to be taken when clearing eligibility to make sure all conditions controlled by all components are met prior to clearing the path to carry traffic.

The current proposal does not introduce a preference between the components acting on this attribute, nor the protocol used to set it. If multiple components are permitted to reset the eligibility flag, the interworking communication between those components to determine if or when eligibility can be restored is out of scope of this document and would be covered on use case document itself.

5. Protocol and model changes

5.1. Active candidate path selection algorithm

As described in Section 4, this proposal introduces a new criteria to the active CP selection process described in section 2.9 of [RFC9256].

5.2. PCEP extensions

PCEP shall be extended to signal the new attribute representing the eligibility of an SR Policy candidate path. A PCE shall be able to change the eligibility status of a delegated LSP and be notified of changes on the eligibility.

5.3. SR policy Yang changes

The eligibility attribute will need to be added to the SR policy candidate path YANG models. NetConf RPC calls can be used to set eligibility of candidate paths to true or false.

5.4. BGP

BGP extensions shall be required to signal and discover the new attribute representing the eligibility of an SR Policy candidate path.

SR Policy CP are sent down via [I-D.ietf-idr-sr-policy-safi] and advertised/published/discovered via BGPLS [I-D.ietf-idr-bgp-ls-sr-policy].

6. IANA considerations

This document includes no request to IANA.

7. Security considerations

This document introduces the concept of eligibility at candidate path construct which is used as an additional criteria during the process of active candidate path selection defined in section 2.9 of [RFC9256]; This document does not expose any additional security challenges to be considered.

Existing security considerations pertaining to SR Policies such as the ones defined in Security Section 10 of [RFC9256] do apply to this document.

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.karboubi-spring-sidlist-optimized-cs-sr]
Karboubi, A., Alaettinoglu, C., Shah, H., Sivalaban, S., and T. Defillipi, "Circuit Style Segment Routing Policies with Optimized SID List Depth, Work in Progress, Internet-Draft,draft-karboubi-spring-sidlist-optimized-cs-sr", 21 February 2025, <<https://datatracker.ietf.org/doc/html/draft-karboubi-spring-sidlist-optimized-cs-sr>>.
- [I-D.ietf-idr-sr-policy-safi]
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-sr-policy-safi-10", 7 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi>>.
- [I-D.ietf-idr-bgp-ls-sr-policy]
Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State, Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-10", 9 December 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-10>>.

Acknowledgements

The authors would like to thank Ketan Talaulikar for his review, comments, and suggestions.

Contributors

Cengiz Alaettinoglu
Ciena
Email: cengiz@ciena.com

Siva Sivabalan
Ciena
Email: ssivabal@ciena.com

Todd Defillipi
Ciena
Email: todd@ciena.com

Ali Zafar
Cisco
Email: zali@cisco.com

Authors' Addresses

Amal Karboubi (editor)
Ciena
Email: akarboub@ciena.com

Himanshu Shah (editor)
Ciena
Email: hshah@ciena.com

Andrew Stone
Nokia
Email: andrew.stone@nokia.com

Christian Schmutzer
Cisco
Email: cschmutz@cisco.com

Praveen Maheshwari
Airtel India
Email: Praveen.Maheshwari@airtel.com