

PCE
Internet-Draft
Intended status: Standards Track
Expires: 4 September 2025

L. Zhang, Ed.
J. Dong
T. Zhou
Huawei
3 March 2025

PCE SR Policy Extensions for Path Scheduling
draft-zzd-pce-sr-policy-scheduling-00

Abstract

Segment Routing (SR) policy enables instantiation of an ordered list of segments with a specific intent for traffic steering. When using SR policy in a time-variant network, delivering the time-variant information associated with paths is necessary in some scenarios.

This document proposes extensions to PCE SR Policy to deliver the schedule information of candidate path (segment list) and its associated attributes.

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 4 September 2025.

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
1.1. Requirements Language	3
2. Motivation	3
2.1. Network with Discontinuous Links	3
2.2. Network with Frequent Topology Changes	4
3. Schedule Information in PCE SR Policy	4
4. Schedule Information TLV	4
4.1. Candidate Paths with Schedule	6
4.2. Segment Lists with Schedule	7
5. Procedures	7
6. Security Considerations	7
7. IANA Considerations	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	9

1. Introduction

[RFC9657] introduces a set of time-variant network use cases where the topology of the network changes predictably. When the networks uses traditional routing protocols, it takes these topology changes as unexpected events and may cause and packet loss. However, the topology changes of these networks can be predicted in advance, therefore some measures can be taken in advance to prevent the packet loss. With this idea, [I-D.ietf-tvr-requirements] describes the requirements of using the time-variant information in a network. In Section 3.4.1 of [I-D.ietf-tvr-requirements], it describes the centralized routing scenarios with time-variant information, in which the network entities receive the time variable information and traffic forwarding rules directly from a logically centralized source(an Orchestrator or network controller). The time-variant information is especially essential when there is a risk that a logically centralized source may loses connectivity with the network entities.

[RFC8664]specifies extensions to the Path Computation Element Communication Protocol (PCEP) that allow a stateful PCE to compute and initiate Traffic-Engineering (TE) paths, as well as a Path Computation Client (PCC) to request a path subject to certain constraints and optimization criteria in SR networks.

[I-D.ietf-pce-segment-routing-policy-cp] extends [RFC8664] to support

signaling SR Policy Candidate Paths as PCEP LSPs and to signal candidate paths of the SR Policy. It signals SR Policy Candidate Paths as PCEP LSPs and signal Candidate Path membership in an SR Policy by means of the Association mechanism. The segment lists of each candidate path and their associated attributes are signaled by the Path Attribute Object defined in [I-D.ietf-pce-multipath]. However, when using SR Policy in a time-variant network, it can't advertise the schedule information associated with paths.

This document proposes extensions to PCE SR Policy to carry the schedule information of candidate paths/segment lists.

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. Motivation

Most of the time-variant network use cases using PCE SR Policy could benefit from this work. In some cases, carrying the time-variant information with SR Policy is essential.

This section describes the cases that requires extending SR Policy with schedule information.

2.1. Network with Discontinuous Links

In some time-variant network cases, the links between the network entities and network controller may very weak or intermittent, this is very typical in Resource Preservation and Dynamic Reachability network[RFC9657]. In these cases, Real-time SR policy advertising (before changes occur) may not be timely. For example, when a link of an old path is about to be disconnected, the network controller is going to advertise a new path to the headend. However, the link between the headend and the network controller is not available. As a result, the new path cannot be advertised in time, causing packet loss.

Therefore, in these cases, once the links between the headend and network controller are available, the controller need to advertise the paths with schedule information for a period in the future to the headend. Then the headend could determine valid paths in the future based on the schedule information of SR policy.

2.2. Network with Frequent Topology Changes

There are also some time-variant network cases that topology changes frequently. This is very typical when the number of network entities is very large (For example, a Dynamic Reachability network with hundreds or thousands of nodes). In this kind of time-variant network, a path from one network entity to another changes frequently, sometimes it can only be maintained for a few minutes or seconds.

Considering that there are multiple paths in a network that computed by the controller, the SR Policies with candidate paths may be advertised to the headend every few seconds. It poses great challenging to the network controller. However, using schedule information could advertise several paths at a time, which greatly mitigate the pressure of network controllers.

3. Schedule Information in PCE SR Policy

4. Schedule Information TLV

[RFC8934] defines the SCHED-LSP-ATTRIBUTE and SCHED-PD-LSP-ATTRIBUTE TLV to indicate the LSP is a non-periodical scheduling LSP or a periodical scheduling LSP. However, it can't express a LSP with complex schedules. On the one hand, the format of these TLVs are very simple, each TLV can only describes one duration or a periodical duration, on the other hand, it requires that only one SCHED-LSP-ATTRIBUTE TLV SHOULD be present in the LSP object, which means each scheduling LSP can only have one duration or periodical duration.

Therefore, the extensions of [RFC8934] could be applicable in some cases with simple schedules, but it is not flexible enough to be applied in the cases with complex schedules (such as the cases listed in Section 2). A more general format of Schedule Information TLV is defined in this draft to cover different kind of cases.

The schedule information TLV indicates one or more valid durations. The format of Schedule Information TLV is shown as follows:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Type										Length																													
Schedules																																							

Figure 1: Schedule Information TLV

Type: TBD1

Length: the size of the value field in octets.

Schedules: one or more schedules, each schedule indicates the duration when the candidate path (segment list) is active. The format of each schedule is shown as follows:

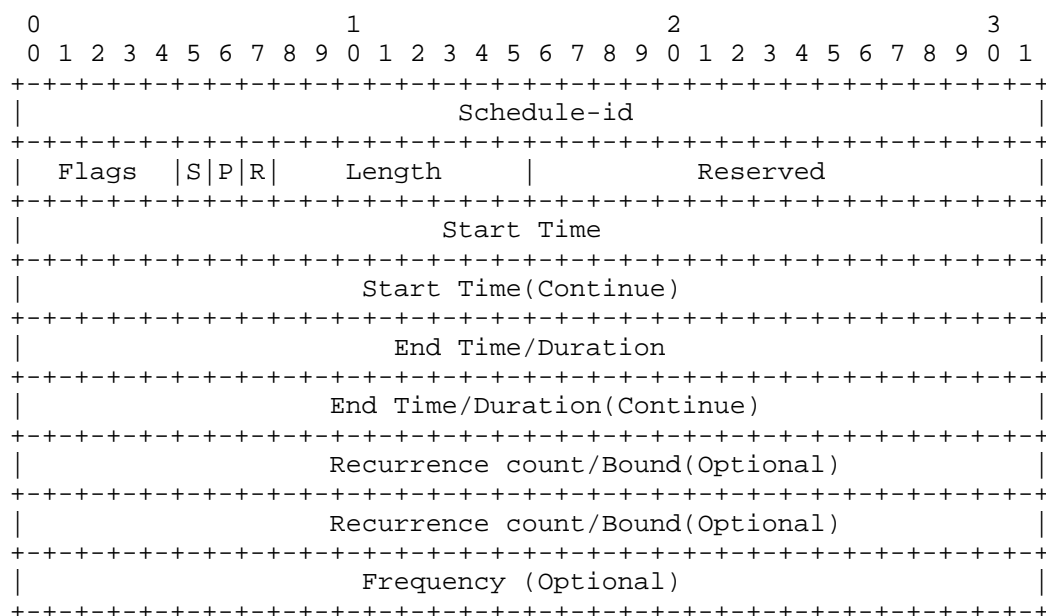


Figure 2: Format of Schedules

Schedule-id: 32-bit value, the unique identifier to distinguish each schedule within a SR Policy, this value is allocated by the SR Policy generator.

Flags: 8 bits, currently only 3 bits are used, the other bits are reserved.

Length: 8 bits, indicates the length of this schedule in octets.

S (Schedule type): one-bit flag to indicate the type of a schedule. If S=0, it indicates the schedule only has one instance, the Recurrence Count/Bound, Frequency and Interval field should not be included in the sub-TLV; If S=1, it indicates the schedule has multiple instances, the Recurrence Count/Bound, Frequency and Interval field should be included.

P (Period type): one-bit flag to indicate the description type of a period. if P=1, then the period is described by a start time field and an end time field; If P =0, then the period is described by a start time field and a duration time field.

R (Recurrence bound type): one-bit flag to indicate the how to determine whether the recurrence is end. if R=1, then the end of recurrence is determined by a detail timepoint; If R = 0, then the end of the recurrence is determined by the number of occurrences.

Start Time: 64-bit value, the number of seconds since the epoch, it indicates when the candidate path (segment list) and its associated attributes start to take effect.

End Time/Duration: 64-bit value, if the flag P=1, then it is the number of seconds since the epoch, it indicates when the candidate path (segment list) and its associated attributes becomes ineffective. If the flag P=0, then it is the number of seconds since the Start Time, it indicates how long the candidate path (segment list) and its associated attributes are effective.

Recurrence Count/Bound(optional): 64-bit value, this field SHOULD be included when the flag P is set to 1. When the flag R=0, then this field indicates the max number of occurrences. For example, if it is set to 2, then the schedule will repeat twice with the specified Frequency and Interval. When the flag R=1, then tis field indicates the bounded timepoint of recurrence, it is described by the number of seconds since the epoch.

Frequency(optional): 32-bit value, this field should be included when the flag S is set to 1. It is the numbers of seconds since the Start Time of an instance to the Start Time of next instance. This field indicates the recurrence frequency for all the instance of this schedule.

4.1. Candidate Paths with Schedule

As described in [I-D.ietf-pce-segment-routing-policy-cpl], SR Policy is represented by a new type of PCEP Association, called the SR Policy Association. The SR Candidate Paths of an SR Policy are represented by the PCEP LSPs within the same SRPA.

When applying schedules to a Candidate Path of an SR Policy, the LSP Object[RFC8231] is required to be extended to support the Schedule Information TLV. The Schedule Information TLV could be an optional TLV present in the LSP Object.

4.2. Segment Lists with Schedule

When there are multiple segment lists within an SR Policy Candidate Paht, [I-D.ietf-pce-multipath] defines the Path Attribute object(PATH-ATTRIB) to carry per-path information. When applying schedules to a Segment List, the PATH-ATTRIB object is required to be extended to support the Schedule Information TLV. The Schedule Information TLV could be an optional TLV present in the LSP Object.

5. Procedures

TBD

6. Security Considerations

TBD

7. IANA Considerations

IANA maintains a sub-registry "PCEP TLV Type Indicators" in the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA is requested to make the following allocations from this sub-registry.

Value	Description	Reference
TBD1	Schedule Information (SI) TLV	This document

Table 1

8. References

8.1. Normative References

[I-D.ietf-tvr-requirements]
King, D., Contreras, L. M., Sipos, B., and L. Zhang, "TVR (Time-Variant Routing) Requirements", Work in Progress, Internet-Draft, draft-ietf-tvr-requirements-04, 13 September 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-tvr-requirements-04>>.

- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [I-D.ietf-pce-segment-routing-policy-cp] Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-22, 25 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-22>>.
- [I-D.ietf-pce-multipath] Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., Peng, S., and G. S. Mishra, "PCEP Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, draft-ietf-pce-multipath-12, 8 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-12>>.
- [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>>.
- [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>>.

8.2. Informative References

- [RFC9657] Birrane, III, E., Kuhn, N., Qu, Y., Taylor, R., and L. Zhang, "Time-Variant Routing (TVR) Use Cases", RFC 9657, DOI 10.17487/RFC9657, October 2024, <<https://www.rfc-editor.org/info/rfc9657>>.
- [RFC8934] Chen, H., Ed., Zhuang, Y., Ed., Wu, Q., and D. Ceccarelli, "PCE Communication Protocol (PCEP) Extensions for Label Switched Path (LSP) Scheduling with Stateful PCE", RFC 8934, DOI 10.17487/RFC8934, October 2020, <<https://www.rfc-editor.org/info/rfc8934>>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.

Authors' Addresses

Li Zhang (editor)
Huawei
Beiqing Road
Beijing
China
Email: zhangli344@huawei.com

Jie Dong
Huawei
Email: jie.dong@huawei.com

Tianran Zhou
Huawei
Email: zhoutianran@huawei.com