

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 13, 2022

A. Tokar
S. Sidor
Cisco Systems, Inc.
S. Sivabalan
Ciena
S. Peng
Huawei Technologies
M. Negi
RtBrick Inc
July 12, 2021

**Carrying SID Algorithm information in PCE-based Networks.
draft-tokar-pce-sid-algo-04**

Abstract

The Algorithm associated with a prefix Segment-ID (SID) defines the path computation Algorithm used by Interior Gateway Protocols (IGPs). This information is available to controllers such as the Path Computation Element (PCE) via topology learning. This document proposes an approach for informing headend routers regarding the Algorithm associated with each prefix SID used in PCE-computed paths, as well as signalling a specific SID algorithm as a constraint to the PCE.

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.

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 January 13, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Object Formats	4
3.1.	OPEN Object	4
3.1.1.	SR PCE Capability Sub-TLV	4
3.1.2.	SRv6 PCE Capability sub-TLV	4
3.2.	SR-ERO Subobject	5
3.3.	SRv6-ERO Subobject	5
3.4.	LSPA Object	5
4.	Operation	6
4.1.	SR-ERO and SRv6-ERO Encoding	6
4.2.	SID Algorithm Constraint	7
5.	Security Considerations	7
6.	IANA Considerations	7
6.1.	SR Capability Flag	7
6.2.	SRv6 PCE Capability Flag	7
6.3.	SR-ERO Flag	8
6.4.	SRv6-ERO Flag	8
6.5.	PCEP TLV Types	8
7.	Normative References	8
Appendix A.	Contributors	9
	Authors' Addresses	9

[1.](#) Introduction

A PCE can compute SR-TE paths using SIDs with different Algorithms depending on the use-case, constraints, etc. While this information is available on the PCE, there is no method of conveying this information to the headend router.

Similarly, the headend can also compute SR-TE paths using different Algorithms, and this information also needs to be conveyed to the PCE for collection or troubleshooting purposes. In addition, in the case of multiple (redundant) PCEs, when the headend receives a path from the primary PCE, it needs to be able to report the complete path information - including the Algorithm - to the backup PCE so that in HA scenarios, the backup PCE can verify the prefix SIDs appropriately.

An operator may also want to constrain the path computed by the PCE to a specific SID Algorithm, for example, in order to only use SID Algorithms for a low-latency path. A new TLV is introduced for this purpose.

Refer to [[RFC8665](#)] and [[RFC8667](#)] for details about the prefix SID Algorithm.

This document is extending:

- o the SR PCE Capability Sub-TLV and the SR-ERO subobject - defined in [[RFC8664](#)]
- o the SRv6 PCE Capability sub-TLV and the SRv6-ERO subobject - defined in [[I-D.ietf-pce-segment-routing-ipv6](#)]

A new TLV for signalling SID Algorithm constraint to the PCE is also introduced, to be carried inside the LSPA object, which is defined in [[RFC5440](#)].

The mechanisms described in this document are equally applicable to both SR-MPLS and SRv6.

2. Terminology

The following terminologies are used in this document:

ERO: Explicit Route Object

IGP: Interior Gateway Protocol

NAI: Node or Adjacency Identifier.

PCE: Path Computation Element

PCEP: Path Computation Element Protocol.

SID: Segment Identifier.

SR: Segment Routing.

SR-TE: Segment Routing Traffic Engineering.

LSP: Label Switched Path.

LSPA: Label Switched Path Attributes.

3. Object Formats

3.1. OPEN Object

3.1.1. SR PCE Capability Sub-TLV

A new flag S is proposed in the SR PCE Capability Sub-TLV introduced in [Section 4.1.2 of \[RFC8664\]](#) in Path Computation Element Communication Protocol (PCEP) to indicate support for SID Algorithm field in the SR-ERO subobject.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|      Type=26                 |      Length=4                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Reserved                 |  Flags |S|N|X|      MSD      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

3.1.2. SRv6 PCE Capability sub-TLV

A new flag S is proposed in the SRv6 PCE Capability sub-TLV introduced in 4.1.1 of [\[I-D.ietf-pce-segment-routing-ipv6\]](#) in Path Computation Element Communication Protocol (PCEP) to indicate support for SID Algorithm field in the SRv6-ERO subobject.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|      Type=TBD1                 |      Length                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Reserved                 |      Flags      |S|N|X|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  MSD-Type   | MSD-Value   |  MSD-Type   |  MSD-Value   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                                     ...                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  MSD-Type   | MSD-Value   |      Padding      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


3.2. SR-ERO Subobject

The SR-ERO subobject encoding is extended with new flag "A" to indicate if the Algorithm field is included after other optional fields.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|  Type=36  |   Length   |  NT  |   Flags  |A|V|F|S|C|M|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SID (optional)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                                     NAI (variable, optional)                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Reserved                                     | Algorithm |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

3.3. SRv6-ERO Subobject

The SRv6-ERO subobject encoding is extended with new flag "A" to indicate if the Algorithm field is included after other optional fields.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|L|  Type=TBD3 |   Length   |  NT  |   Flags  |A|V|T|F|S|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Reserved                                     | Endpoint Behavior |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SRv6 SID (optional)                                     |
|                                     (128-bit)                                     |
|                                     |                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
//                                     NAI (variable, optional)                                     //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     SID Structure (optional)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Reserved                                     | Algorithm |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

3.4. LSPA Object

A new TLV for the LSPA Object with TLV type=TBD3 is introduced to carry the SID Algorithm constraint. This TLV SHOULD only be used when PST (Path Setup type) = SR or SRv6.

The format of the SID Algorithm TLV is as follows:




Figure 1: SID Algorithm TLV Format

The code point for the TLV type is TBD3. The TLV length is 4 octets.

The 32-bit value is formatted as follows.

Reserved: MUST be set to zero by the sender and MUST be ignored by the receiver.

Flags: This document defines the following flag bits. The other bits MUST be set to zero by the sender and MUST be ignored by the receiver.

- * L (Loose): If set to 1, the PCE MAY insert SIDs with a different Algorithm, but it MUST prefer the specified Algorithm whenever possible.

Algorithm: SID Algorithm the PCE MUST take into account while computing a path for the LSP.

4. Operation

4.1. SR-ERO and SRv6-ERO Encoding

PCEP speaker MAY set the A flag and include the Algorithm field in SR-ERO or SRv6-ERO subobject if the S flag was advertised by both PCEP speakers.

If PCEP peer receives SR-ERO subobject with the A flag set or with the SID Algorithm included, but the S flag was not advertised, then such PCEP message must be rejected with PCError as described in [Section 7.2 of \[RFC5440\]](#)

The Algorithm field MUST be included after optional SID, NAI or SID structure and length of SR-ERO or SRv6-ERO subobject MUST be increased with additional 4 bytes for Reserved and Algorithm field.

4.2. SID Algorithm Constraint

In order to signal a specific SID Algorithm constraint to the PCE, the headend MUST encode the SID ALGORITHM TLV inside the LSPA object.

When the PCE receives a SID Algorithm constraint, it MUST only take prefix SIDs with the specified Algorithm into account during path computation. However, if the L flag is set in the SID Algorithm TLV, the PCE MAY insert prefix SIDs with a different Algorithm in order to successfully compute a path.

If the PCE is unable to find a path with the given SID Algorithm constraint, it MUST bring the LSP down.

SID Algorithm does not replace the Objective Function defined in [RFC5541]. The SID Algorithm constraint acts as a filter, restricting which SIDs may be used as a result of the path computation function.

5. Security Considerations

No additional security measure is required.

6. IANA Considerations

6.1. SR Capability Flag

IANA maintains a sub-registry, named "SR Capability Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flags field of the SR-PCE-CAPABILITY TLV. IANA is requested to make the following assignment:

Value	Description	Reference
TBD1	SID Algorithm Capability	This document

6.2. SRv6 PCE Capability Flag

IANA was requested in [I-D.ietf-pce-segment-routing-ipv6] to create a sub-registry, named "SRv6 PCE Capability Flags", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flags field of SRv6-PCE-CAPABILITY sub-TLV. IANA is requested to make the following assignment:

Value	Description	Reference
TBD2	SID Algorithm Capability	This document

6.3. SR-ERO Flag

IANA maintains a sub-registry, named "SR-ERO Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flags field of the SR-ERO Subobject. IANA is requested to make the following assignment:

Value	Description	Reference
TBD3	SID Algorithm Flag	This document

6.4. SRv6-ERO Flag

IANA was requested in [[I-D.ietf-pce-segment-routing-ipv6](#)], named "SRv6-ERO Flag Field", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the Flags field of the SRv6-ERO subobject. IANA is requested to make the following assignment:

Value	Description	Reference
TBD4	SID Algorithm Flag	This document

6.5. PCEP TLV Types

IANA is requested to allocate a new TLV type for the new LSPA TLV specified in this document.

Value	Description	Reference
TBD5	SID Algorithm	This document

7. Normative References

- [I-D.ietf-pce-segment-routing-ipv6]
Li, C., Negi, M., Sivabalan, S., Koldychev, M., Kaladharan, P., and Y. Zhu, "PCEP Extensions for Segment Routing leveraging the IPv6 data plane", [draft-ietf-pce-segment-routing-ipv6-09](#) (work in progress), May 2021.
- [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>>.

- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", [RFC 5541](#), DOI 10.17487/RFC5541, June 2009, <<https://www.rfc-editor.org/info/rfc5541>>.
- [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., 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>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [RFC 8665](#), DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [RFC 8667](#), DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.

[Appendix A](#). Contributors

Mike Koldychev
Cisco Systems
Kanata, Ontario
Canada

Email: mkoldych@cisco.com

Authors' Addresses

Alex Tokar
Cisco Systems, Inc.
Eurovea Central 3.
Pribinova 10
Bratislava 811 09
Slovakia

Email: atokar@cisco.com

Samuel Sidor
Cisco Systems, Inc.
Eurovea Central 3.
Pribinova 10
Bratislava 811 09
Slovakia

Email: ssidor@cisco.com

Siva Sivabalan
Ciena
385 Terry Fox Drive
Kanata, Ontario K2K 0L1
Canada

Email: msiva282@gmail.com

Shuping Peng
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: pengshuping@huawei.com

Mahendra Singh Negi
RtBrick Inc
Bangalore, Karnataka
India

Email: mahend.ietf@gmail.com

