idr
Internet-Draft
Intended status: Standards Track
Expires: 3 August 2026

Q. Xiong
ZTE Corporation
P. Liu
China Mobile
R. Gandhi
Cisco Systems, Inc.

30 January 2026

PCEP Extension for Bounded Latency
draft-xiong-pce-detnet-bounded-latency-07

Abstract

In certain networks, such as Deterministic Networking (DetNet), it is required to consider the bounded latency for path selection. This document describes the extensions for Path Computation Element Communication Protocol (PCEP) to carry the bounded latency constraints and distribute deterministic paths for end-to-end path computation in deterministic services.

Table of Contents

1. Introduction 2
1.1. Requirements Language 4
2. Terminology 4
3. PCEP Extensions 4
3.1. METRIC Object 4
3.1.1. End-to-End Minimum Latency Metric 4
3.1.2. End-to-End Maximum Latency Metric 4
3.1.3. End-to-End Latency Variation Metric 5
3.2. LSP Object 5
3.3. Deterministic Path ERO Subobject 5
3.4. Deterministic Path RRO Subobject 7
4. Operations 7
4.1. Calculation of End-to-end Bounded Latency 7
4.2. Metric types 8
4.3. ERO and RRO Subobjects 8
5. Security Considerations 9
6. IANA Considerations 10
6.1. New Metric Types 10
6.2. New LSP-EXTENDED-FLAG Flag Registry 10
6.3. New ERO and RRO Subobject 10
7. Acknowledgements 10
8. References 10
8.1. Normative References 10
8.2. Informative References 13
Authors' Addresses 13

1. Introduction

[RFC5440] describes the Path Computation Element Protocol (PCEP) which is used between a Path Computation Element (PCE) and a Path Computation Client (PCC) (or other PCE) to enable computation of Multi-protocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP). PCEP Extensions for the Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of MPLS-TE and Generalized MPLS (GMPLS) tunnels.
As depicted in [RFC4655], a PCE MUST be able to compute the path of a TE LSP by operating on the TED and considering bandwidth and other constraints applicable to the TE LSP service request. The constraint parameters are provided such as metric, bandwidth, delay, affinity, etc. However these parameters did not take into account the bounded latency requirements.

According to [RFC8655]}, Deterministic Networking (DetNet) operates at the IP layer and delivers service which provides extremely low data loss rates and bounded latency within a network domain. The bounded latency indicates the minimum and maximum end-to-end latency from source to destination and bounded jitter (packet delay variation). [I-D.ietf-detnet-scaling-requirements] has described the enhanced requirements for DetNet data plane including the information used by functions ensuring deterministic latency should be supported. And queuing mechanisms and solutions require different information to help the functions of ensuring deterministic latency, including regulation, queue management. [I-D.ietf-detnet-dataplane-taxonomy] has defined the classification criteria and the suitable categories for this solutions.

The computing method of end-to-end delay bounds is defined in [RFC9320]. It is the sum of the 6 delays in DetNet bounded latency model. And these delays should be measured and collected by IGP, but the related mechanisms are out of this document. The end-to-end delay bounds can also be computed as the sum of non queuing delay bound and queuing delay bound along the path. The upper bounds of non queuing delay are constant and depend on the specific network and the value of queuing delay bound depends on the queuing mechanisms deployed along the path. The queuing delay may differ notably in their specific queuing solutions, which should be selected and calculated by the controller (or PCE). The deterministic latency information related to each queuing mechanism should also be distributed.

As per [I-D.ietf-detnet-controller-plane-framework], explicit path should be calculated and established in control plane to guarantee the deterministic transmission. The corresponding IS-IS and OSPF extensions are specified in [I-D.peng-lsr-deterministic-traffic-engineering]. When the PCE is deployed, the path computation should be applicable for deterministic networks. It is required that bounded latency including minimum and maximum end-to-end latency and bounded delay variation are considered during the deterministic path selection for PCE. The bounded latency constraints should be extended for PCEP. Moreover, the queuing-based parameters along the deterministic path should be provided to the PCC after the path computation such as deterministic latency information.

This document describes the extensions for PCEP to carry bounded latency constraints and distribute deterministic paths for end-to-end path computation in deterministic services.

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 terminology is defined as [RFC8655] and [RFC5440].

3. PCEP Extensions

3.1. METRIC Object

The METRIC object is defined in Section 7.8 of [RFC5440], comprising metric-value and metric-type (T field), and a flags field, comprising a number of bit flags (B bit and C bit). This document defines three types for the METRIC object to represent the end-to-end bounded latency.

3.1.1. End-to-End Minimum Latency Metric

This document proposes the end-to-end minimum latency metric in PCEP to represent the lower bound of the end-to-end delay. The extensions for End-to-End Minimum Latency Metric are as following shown:

*T=TBD1: End-to-End Minimum Latency Metric.

*The value of End-to-End Minimum Latency Metric is the encoding in units of microseconds with 32 bits.

*The B bit MUST be set to suggest a minimum bound for the end-to-end delay of deterministic path. The end-to-end delay must be no less than or equal to the value.

3.1.2. End-to-End Maximum Latency Metric

This document proposes the end-to-end maximum latency metric in PCEP to represent the upper bound of the end-to-end delay. The extensions for End-to-End Maximum Latency Metric are as following shown:

*T=TBD2: End-to-End Maximum Latency Metric.

*The value of End-to-End Maximum Latency Metric is the encoding in units of microseconds with 32 bits.

*The B bit MUST be set to suggest a maximum bound for the end-to-end delay of deterministic path. The end-to-end delay must be less than or equal to the value.

3.1.3. End-to-End Latency Variation Metric

This document proposes the end-to-end latency variation metric in PCEP to represent the difference between the end-to-end upper latency and the end-to-end lower latency along a deterministic path. The extensions for End-to-End Latency Variation Metric are as following shown:

*T=TBD3: End-to-End Latency Variation Metric.

*The value of End-to-End Latency Variation Metric is the encoding in units of microseconds with 32 bits.

*The B bit MUST be set to suggest a maximum bound for the end-to-end latency variation of deterministic path. The end-to-end latency variation must be less than or equal to the value.

3.2. LSP Object

The LSP Object is defined in Section 7.3 of [RFC8231]. This document defines a new flag (D-flag) to present the deterministic path for the LSP-EXTENDED-FLAG TLV carried in LSP Object as defined in [RFC9357].

D (Request for Deterministic Path) : If the bit is set to 1, it indicates that the PCC requests PCE to compute the deterministic path. A PCE would also set this bit to 1 to indicate that the deterministic path is included by PCE and encoded in the PCRep, PCUpd or PCInitiate message.

3.3. Deterministic Path ERO Subobject

The ERO (Explicit Route Object) specified in [RFC3209] and [RFC5440] can be used to carry a set of computed paths. In order to carry deterministic latency information, this document defines a new optional ERO subobject referred to as the Deterministic Path ERO subobject (DP-ERO). An ERO carrying a deterministic path consists of one or more ERO subobjects, and it MUST carry DP-ERO subobjects.

An DP-ERO subobject is formatted as shown in the following figure:

```
       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=TBD4  |     Length    |     Class     |    DLI  Type  |
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
      //  Deterministic Latency Information(variable, optional)       //
      +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1 DP-ERO subobject

where:

*L (1bit): The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the explicit route. If the bit is not set, the subobject represents a strict hop in the explicit route.

*Type (8bits): Set to TBD4.

*Length (8bits): Contains the total length of the subobject in octets.
The Length MUST be at least 8 and MUST be a multiple of 4.

*Class (8bits): indicates the deterministic forwarding class.

*Deterministic Latency Information(DLI) Type (8bits): indicates the type of deterministic latency information with related queuing and scheduling metadata and it aglined with the suitable categories as defined in [I-D.ietf-detnet-dataplane-taxonomy] and shown in Figure 2.

Value	DLI Type
0x0000	Unassigned
0x0001	Right-bounded
0x0002	Flow level periodic bounded
0x0003	Class level periodic bounded
0x0004	Flow level non-periodic bounded
0x0005	Class level non-periodic bounded
0x0006	Flow level rate based unbounded
0x0007	Flow level rate based left-bounded

Figure 2 DLI Type

*Deterministic Latency Information(DLI) (variable): indicates the corresponding deterministic latency parameters. The encoding format of each DLI depends on the value of the DLI type and aligns with the deterministic latency information as defined in section 4.2 [I-D.xiong-detnet-data-fields-edp].

3.4. Deterministic Path RRO Subobject

The Deterministic Path RECORD_ROUTE Object (DP-RRO) subobject is OPTIONAL. If used, it is carried in the RECORD_ROUTE Object (RRO). The subobject uses the standard format of an RRO subobject. The format of the DP-RRO subobject is the same as that of the DP-ERO subobject, but without the L flag.

4. Operations

4.1. Calculation of End-to-end Bounded Latency

As per [RFC9320], the end-to-end delay bound can be computed as the sum of Output delay, Link delay, Frame preemption delay, Processing delay, Regulation delay and Queuing delay along a deterministic path like following:

*per-hop_delay_bound = sum{Output delay + Link delay + Frame preemption delay + Processing delay + Regulation delay + Queuing delay}.

*end_to_end_delay_bound = sum{per-hop_delay_bound(h),(h=1,2,...H)}.

As per [RFC9320], it also can be encoded as the sum of non queuing delay bound and queuing delay bound along the deterministic path. Per-hop non queuing delay bound is the sum of the bounds over delays including Output delay, Link delay, Frame preemption delay and Processing delay and per-hop queuing delay bound is the sum of Regulation delay and Queuing delay like following:

*end_to_end_delay_bound = non_queuing_delay_bound + queuing_delay_bound.

As per [RFC9320], the end-to-end delay variation can be encoded as the sum of non queuing delay variation and queuing delay variation along the deterministic path like following:

*end_to_end_delay_variation = non_queuing_delay_variation + queuing_delay_variation.

Moreover, as discussed in [I-D.ietf-detnet-dataplane-taxonomy], the end-to-end bounded latency calculation includes the bounded delay and variation. The calculation of end-to-end bounded delay and variation will differ in each queuing solution. For example, the end-to-end delay variation is 2 times of the cycle ID when selecting cyclic-based queuing mechanism.

4.2. Metric types

The PCE needs to collect the value of the delays as per [RFC9320] and related parameters by IGP, calculate the bounded latency, select a deterministic path with a specific queuing mechanism which meet the requirements and configure the related parameters to a PCC. The PCC MAY use the end-to-end bounded latency metrics in a Path Computation Request (PCReq) message to request a deterministic path meeting the end-to-end bounded latency requirements. A PCE MAY use the metrics in a Path Computation Reply (PCRep) message along with a NO-PATH object in the case where the PCE cannot compute a path meeting this constraints. A PCE can also use the metrics to send the computed end-to-end bounded latency to the PCC.

4.3. ERO and RRO Subobjects

A PCC can request the computation of deterministic path and a PCE may respond with PCRep message. And the deterministic path can also be initiated by PCE with PCInitiate or PCUpd message in stateful PCE mode. When the D bit in LSP object is set to 1 within the message, it indicates to request the calculation of deterministic path. When the bit is set in Metric object to indicate the end-to-end bounded latnecy metric, the PCE should calculate the end-to-end latency bound to select the optimal deterministic path to meet the requirements.

The DP-ERO subobject can be carried along the path to indicate the deterministic path and related information. The deterministic path being received by PCC encoded in DP-ERO, which carry the deterministic latency information. And the PCC may insert the deterministic latency information as the DetNet-specific metadata into the packet headers to achieve the deterministic forwarding.

The set of computed paths can be specified by means of ERO [RFC3209], SR-RRO [RFC8664] and SRv6-ERO [RFC9603] subobjects. When the D bit in LSP object is set to 1, a DP-ERO subobject which carrying the deterministic path information MAY be inserted directly after the existing identifying subobjects such as ERO [RFC3209] , SR-ERO [RFC8664] and SRv6-ERO [RFC9603]. The PCE will select only one DLI type from seven categories and compute the deterministic path with different DLI in each node along the path.

A DP-ERO subobject corresponds to be a preceding subobject which can not be the first subobject. The PCE will select one DLI type from seven categories and compute the deterministic path with different DLI in each node along the path. For example, when the result is A, B, C in SR networks and the results with deterministic path will be like following:

```
Deterministic path example with DLI type is 2:
SR-ERO subobject (SID A) + DP-ERO subobject (DLI type = 2, DLI carry Flow level periodic
bounded info at node A) ->
SR-ERO subobject (SID B) + DP-ERO subobject (DLI type = 2, DLI carry Flow level periodic
bounded info at node B) ->
SR ERO subobject (SID C) + DP-ERO subobject (DLI type = 2, DLI carry Flow level periodic
bounded info at node C)

Deterministic path example with DLI type is 3:
SR-ERO subobject (SID A) + DP-ERO subobject (DLI type = 3, DLI carry Class level periodic
 bounded info at node A) ->
SR-ERO subobject (SID B) + DP-ERO subobject (DLI type = 3, DLI carry Class level periodic
 bounded info at node B) ->
SR-ERO subobject (SID C) + DP-ERO subobject (DLI type = 3, DLI carry Class level periodic
 bounded info at node C)
```

The DP-RRO subobject can be also carried directly after the existing identifying RRO subobjects such as RRO [RFC3209], SR-RRO [RFC8664] and SRv6-RRO [RFC9603].

5. Security Considerations

Security considerations for DetNet are covered in the DetNet architecture [RFC8655], DetNet security considerations [RFC9055] and DetNet control plane [I-D.ietf-detnet-controller-plane-framework]. This document defines a new D bit and DP-ERO subobject for deterministic path in PCEP, which do not introduce any new security considerations beyond those already listed in [RFC5440],[RFC8231] and

[RFC9357].

6. IANA Considerations

6.1. New Metric Types

This document defines two new metric type for the PCEP. IANA is requested to allocate the following codepoint in the PCEP "METRIC Object T Field" registry:

Value	Description	Reference
TBD1	End-to-End Minimum Latency Metric	This document
TBD2	End-to-End Maximum Latency Metric	This document
TBD3	End-to-End Latency Variation Metric	This document

6.2. New LSP-EXTENDED-FLAG Flag Registry

[RFC9357] defines the LSP-EXTENDED-FLAG TLV. IANA is requested to make allocations from the Flag field registry, as follows:

Bit	Description	Reference
D flag	Request for Deterministic Path	This document

6.3. New ERO and RRO Subobject

This document defines a new subobject type for the PCEP explicit route object (ERO). The code points for subobject types of these objects is maintained in the RSVP parameters registry, under the EXPLICIT_ROUTE and RECORD_ROUTE objects. IANA is requested to confirm the following allocations in the RSVP Parameters registry for each of the new subobject types defined in this document.

Object	Subobject	Subobject Type
EXPLICIT_ROUTE	DP-ERO (PCEP-specific)	TBD4
RECORD_ROUTE	DP-RRO (PCEP-specific)	TBD4

7. Acknowledgements

The authors would like to thank Dhruv Dhody, Andrew Stone, Lou Berger, Janos Farkas for their review, suggestions and comments to this document.

8. References

8.1. Normative References

[I-D.ietf-detnet-controller-plane-framework]
Malis, A. G., Geng, X., Chen, M., Varga, B., and C. J. Bernardos, "A Framework for Deterministic Networking (DetNet) Controller Plane", Work in Progress, Internet-Draft, draft-ietf-detnet-controller-plane-framework-15, 24 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-controller-plane-framework-15>.

[I-D.ietf-detnet-dataplane-taxonomy]
Joung, J., Geng, X., Peng, S., and T. T. Eckert, "Dataplane Enhancement Taxonomy", Work in Progress, Internet-Draft, draft-ietf-detnet-dataplane-taxonomy-05, 8 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-dataplane-taxonomy-05>.

[I-D.ietf-detnet-scaling-requirements]
Liu, P., Li, Y., Eckert, T. T., Xiong, Q., Ryoo, J., zhushiyin, and X. Geng, "Requirements for Scaling Deterministic Networks", Work in Progress, Internet-Draft, draft-ietf-detnet-scaling-requirements-09, 7 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-detnet-scaling-requirements-09>.

[I-D.xiong-detnet-data-fields-edp]
Xiong, Q., Liu, A., Gandhi, R., and D. Yang, "Data Fields for DetNet Enhanced Data Plane", Work in Progress, Internet-Draft, draft-xiong-detnet-data-fields-edp-03, 1 July 2025, <https://datatracker.ietf.org/doc/html/draft-xiong-detnet-data-fields-edp-03>.

[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>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/rfc/rfc3209>.

[RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/rfc/rfc4655>.

[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/rfc/rfc5440>.

[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>.

[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/rfc/rfc8231>.

[RFC8233] Dhody, D., Wu, Q., Manral, V., Ali, Z., and K. Kumaki, "Extensions to the Path Computation Element Communication Protocol (PCEP) to Compute Service-Aware Label Switched Paths (LSPs)", RFC 8233, DOI 10.17487/RFC8233, September 2017, <https://www.rfc-editor.org/rfc/rfc8233>.

[RFC8655] Finn, N., Thubert, P., Varga, B., and J. Farkas, "Deterministic Networking Architecture", RFC 8655, DOI 10.17487/RFC8655, October 2019, <https://www.rfc-editor.org/rfc/rfc8655>.

[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/rfc/rfc8664>.

[RFC9320] Finn, N., Le Boudec, J.-Y., Mohammadpour, E., Zhang, J., and B. Varga, "Deterministic Networking (DetNet) Bounded Latency", RFC 9320, DOI 10.17487/RFC9320, November 2022, <https://www.rfc-editor.org/rfc/rfc9320>.

[RFC9357] Xiong, Q., "Label Switched Path (LSP) Object Flag Extension for Stateful PCE", RFC 9357, DOI 10.17487/RFC9357, February 2023, <https://www.rfc-editor.org/rfc/rfc9357>.

[RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <https://www.rfc-editor.org/rfc/rfc9603>.

8.2. Informative References

[I-D.peng-lsr-deterministic-traffic-engineering]
Peng, S., "IGP Extensions for Deterministic Traffic Engineering", Work in Progress, Internet-Draft, draft-peng-lsr-deterministic-traffic-engineering-04, 8 January 2026, <https://datatracker.ietf.org/doc/html/draft-peng-lsr-deterministic-traffic-engineering-04>.

[RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/rfc/rfc9055>.

Authors' Addresses

Quan Xiong
ZTE Corporation
Email: xiong.quan@zte.com.cn

Peng Liu
China Mobile
Email: liupengyjy@chinamobile.com

Rakesh Gandhi
Cisco Systems, Inc.
Email: rgandhi@cisco.com