

Workgroup: Network Working Group
Internet-Draft:
draft-zhang-pce-enhanced-detnet-03
Published: 10 July 2023
Intended Status: Standards Track
Expires: 11 January 2024
Authors: L. Zhang X. Geng T. Zhou
 Huawei Huawei Huawei

PCEP for Enhanced DetNet

Abstract

PCEP is used to provide a communication between a PCC and a PCE. This document defines the extensions to PCEP to support the bounded-latency path computation. Specifically, two new objects and three new TLVs are defined for the transmission of bounded latency information between PCC and PCE to guarantee the bounded latency transmission in control plane.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 11 January 2024.

Copyright Notice

Copyright (c) 2023 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. Terminology](#)
- [3. Object Formats](#)
 - [3.1. Open Object](#)
 - [3.1.1. Bounded Latency Capability TLV](#)
 - [3.2. RP Object](#)
 - [3.2.1. BLI Type TLV](#)
 - [3.3. Traffic Model Object](#)
 - [3.4. BLI Object](#)
 - [3.4.1. BLI List TLV](#)
 - [3.4.2. Shared BLI TLV](#)
- [4. SR Policy for BLI](#)
- [5. IANA Considerations](#)
 - [5.1. New TLV Type](#)
 - [5.2. New Object](#)
- [6. Security Considerations](#)
- [7. Acknowledgements](#)
- [8. Normative References](#)
- [Authors' Addresses](#)

1. Introduction

[[RFC8665](#)] provides the overall architecture for Deterministic Networking (DetNet), which provides the capability to carry specified unicast or multicast data flows with extremely low data loss rates and bounded end-to-end latency within a network domain. Based on this, [[draft-finn-detnet-bounded-latency](#)] proposed a timing model for sources, destinations, and DetNet transit nodes. Using the model, it provides a methodology to compute end-to-end latency and backlog bounds for various queuing methods.

[[RFC5440](#)] describes the Path Computation Element Protocol (PCEP) for communications between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between two PCEs. PCEP defines the interaction and data format of path calculation requests and path computation replies between PCC and PCE. [[RFC8231](#)] specifies extension to PCEP to enable stateful control of LSPs within and across PCEP sessions in compliance with [[RFC4657](#)].

[[I-D.yzz-detnet-enhanced-data-plane](#)] enhances the DetNet data plane

by introducing Bounded Latency Information (BLI) which facilitates DetNet transit nodes to guarantee the bounded latency transmission in data plane. Based on that, [[I-D.geng-spring-sr-enhanced-detnet](#)] defines how to leverage Segment Routing (SR) and Segment Routing over IPv6 (SRv6) to implement bounded latency.

When a PCE is used to compute paths using PCEP, it is important that the PCE understands the bounded latency requirement and the head end of the path also need to understand the bounded latency information associated with the candidate path.

This document defines the extensions to PCEP to support the bounded-latency path computation. Specifically, two new objects and three new TLVs are defined for the transmission of bounded latency information between PCC and PCE to guarantee the bounded latency transmission in control plane.

2. Terminology

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 [[RFC2119](#)].

3. Object Formats


3.1. Open Object

[[RFC5440](#)] defines the Open object in open message used to specify the PCEP version, Keepalive frequency, DeadTimer, PCEP session ID, and other communication parameters. The Open object may also contain a set of TLVs used to convey various session characteristics.

3.1.1. Bounded Latency Capability TLV

During the PCEP initialization phase, PCEP speakers SHOULD advertise their support of Bounded Latency features, for this reason this document defines the Bounded Latency capability TLV.

A PCEP speaker includes the Bounded Latency capability TLV in the Open object to advertise its support for Bounded Latency features. The format of the Bounded Latency capability TLV is formatted as follows:



Type: To be assigned by IANA.

Length: 16 bits value to indicate the length of the value portion in bytes.

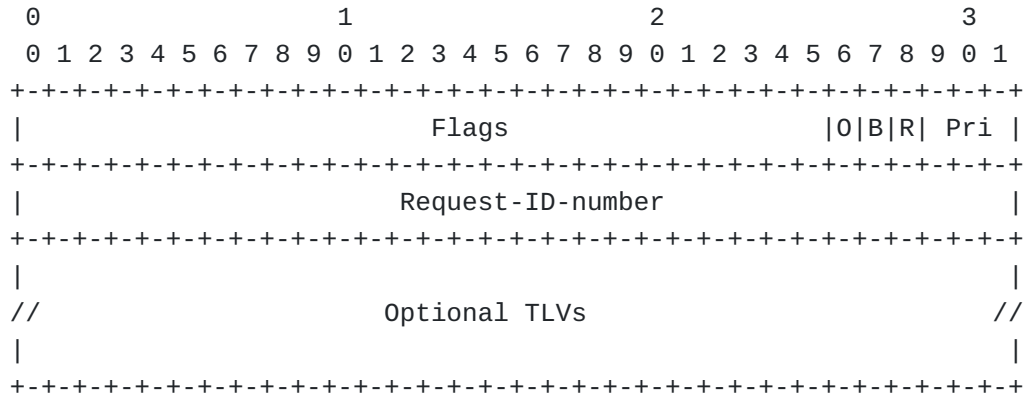
Type-Flag: 16 bits of flags to indicate which kind of BLI Type the speaker supports. A new registry "Bounded Latency Type Flags" is expected to be created. Table 1 shows the assignment of Bounded Latency Type Flags. The speaker sets the defined bit in flag to indicate that it supports this Type of BLI. The undefined bits MUST be set to zero by the sender and MUST be ignored by the receiver.

Bit	BLI Type	
0	Reserved	
1	Time resource ID	32-bit
2	Priority	8-bit
3	End-to-end delay budget	32-bit
4	Local delay budget	32-bit
5	Reserved	
6	Reserved	
7	End-to-end delay variation budget	16-bit
8	Local delay variation budget	16-bit
9-15	Undefined	

Table1: The BLI type and format of each bit flag

3.2. RP Object


The RP (Request Parameters) object is defined in[RFC5440], used to specify various characteristics of the path computation request and MUST be carried within each PCReq and PCRep messages. The format of RP object is as follows:



The detail information about the fields in the RP object is defined in section 7.4 of[RFC5440].

3.2.1. BLI Type TLV

In order to specify the type and format of the BLI associated with candidate path, this document defines a new TLV named BLI type TLV. The BLI type TLV is formatted as follow:



Where:

Type: to be assigned by IANA.

Length: 16 bits value to indicate the length of the value portion in bytes. The value of this field is 4.

BLI Type: 8 bits value to indicate the type of BLI that PCC desires. Table 3 shows the values and their corresponding BLI types.

BLI Type Value	Bounded Latency Information	
0	Reserved	
1	Time resource ID	32-bit
2	Priority	8-bit
3	End-to-end delay variation budget	16-bit
4	Local delay budget	32-bit
5	End-to-end queue delay budget	32-bit
6	Local queue delay budget	32-bit

Table2: The BLI type and format of each value

When PCC needs to request a bounded-latency path, it MUST include the BLI Type TLV in the RP object in PCReq message. If a PCC includes an BLI Type TLV on a path calculation request, then the PCE will reply the specific type of BLI associated with computed path.

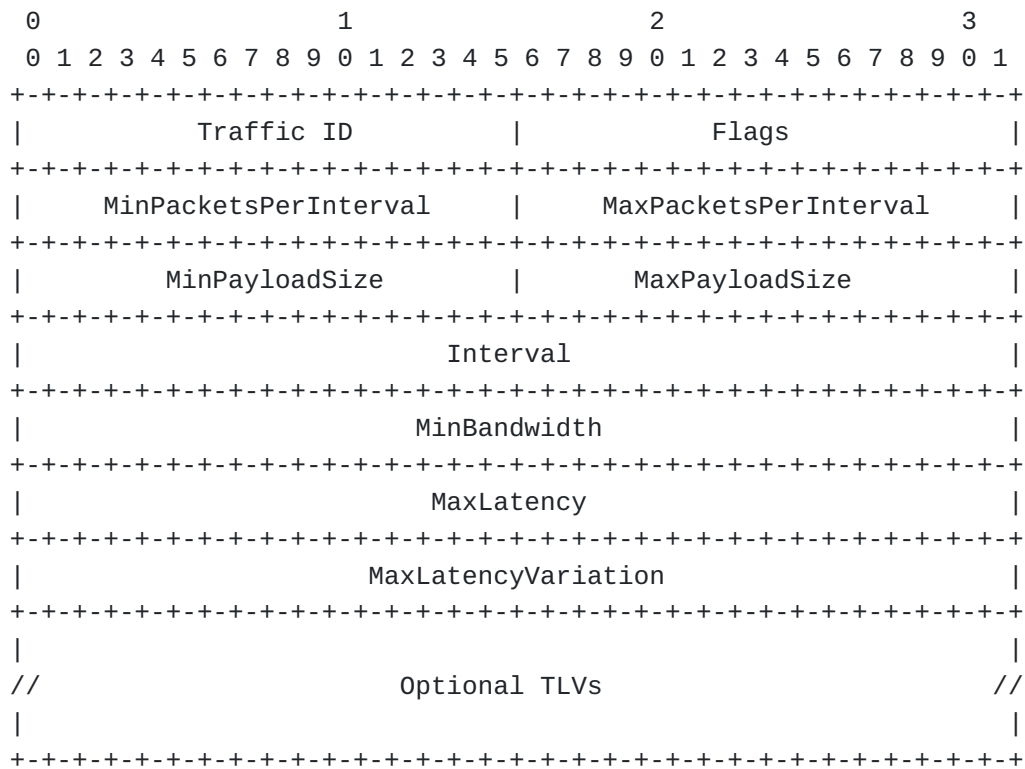
3.3. Traffic Model Object

The Traffic Model Object is optional in the PCReq message and used to specify the traffic model for the bounded-latency path computation. The traffic model object contains a set of fields used to specify the traffic features. [\[RFC9016\]](#) defines the traffic specification of the DetNet flow, which includes a set of attributes to specify how the DetNet Ingress transmits packets for the DetNet flow. Based on that, this document proposes the Traffic Model Object to describe the DetNet flow for bounded-latency path computation.

Traffic Model Object-Class is TBD3;

Traffic Model Object-Type is 1.

The format of the Traffic Model Object is shown in below:



Where:

Traffic ID: The only identification of the specify traffic in PCC. When the PCC need to request a path computation for a traffic, it MUST assign a 16-bit traffic identifier to specify the traffic in local.

Flags: 16 bits of flags. A new registry "Traffic Model Flags" is expected to be created. At the writing time, all flags are unused and undefined.

MinPacketsPerInterval: the minimum number of packets that the Ingress will transmit in one Interval.

MaxPacketsPerInterval: the maximum number of packets that the Ingress will transmit in one Interval.

MinPayloadSize: the minimum payload size that the Ingress will transmit.

MaxPayloadSize: the maximum payload size that the Ingress will transmit.

Interval: the period of time in which the traffic specification is specified.

MinBandwith: the minimum bandwidth that has to be guaranteed for the DetNet traffic.

MaxLatency: the end-to-end maximum latency for a single packet of the DetNet traffic.

MaxLatencyVariation: the difference between the minimum and the maximum end-to-end, one-way latency.

The Traffic Model object body has a variable length and may contain TLVs for the additional attributes of the traffic model. At the writing time there is no TLV defined for Traffic Mode Object.

3.4. BLI Object


In order to support the bounded-latency path computation, a new kind of object named BLI object is defined in this document to indicate the bounded latency information of a candidate path.

The BLI object is optionally carried within a PCRep message so as to indicate the requirement and resource allocation for the candidate path. When a PCC request a bounded-latency path computation and the PCE find out a path satisfying the set of constraints, the PCE MUST include the BLI object in PCRep message.

BLI Object-Class is TBD4.

BLI Object-Type is 1.

The format of BLI object is as follows:



The BLI object body has a variable length and may contain TLVs for the different kinds of BLI. This document defines two kinds of BLI TLV for different scenarios.

3.4.1. BLI List TLV

When all of the nodes in the Explicit Route Object (ERO)[[RFC5440](#)] request different BLI to guarantee bounded latency, a BLI list TLV is defined.

The BLI list sub-TLV is formatted as follows.



Where:

Type: to be assigned by IANA.

Length: 16 bits length value to indicate the length of BLI list in octet.


BLI List [1... m]: 32 bits length bounded latency information, representing the nth BLI in the BLI list.

The BLI in the BLI List corresponds to the node in the ERO object one by one. The length of the BLI List depends on the number of nodes in the ERO object.

3.4.2. Shared BLI TLV

When all of the nodes in the ERO indicated by the sub-object list request BLI to guarantee bounded latency with the same BLI value, the Shared BLI TLV is defined.

The Shared BLI TLV is defined as follows:



Where:

Type: to be assigned by IANA.

Length: 16 bits value to indicate the length of BLI in octet.

BLI: 32 bits value of Bounded Latency Information to guarantee the bounded latency.

4. SR Policy for BLI

[[I-D.ietf-pce-segment-routing-policy-cp](#)] proposes extension to PCEP to support association among candidate paths of a given SR policy. For the bounded latency path, the additional bounded latency information associated with the candidate path SHOULD be carried with SR Policy. Therefore, the additional BLI TLV SHOULD be defined to indicate the bounded-latency requirement and resources allocation for the nodes along the candidate path. For different scenario, different BLI TLV need to be carried by SR policy.

When all of the nodes/adjacencies in the explicit path indicated by the segment list request different BLI to guarantee bounded latency, a BLI list TLV is need to be carried by SR Policy. The BLI list TLV is defined in section 3.4.1.

When all of the nodes/adjacencies in the explicit path indicated by the segment list request BLI to guarantee bounded latency with the same BLI value, a Shared BLI TLV is need to be carried by SR Policy. The Shared BLI TLV is defined in section 3.4.2.

5. IANA Considerations

This document defines four new TLVs and two new Object.

5.1. New TLV Type

Value	Name	Reference
TBD1	Bounded-Latency Capability TLV	This document
TBD2	BLI Type TLV	This document
TBD5	BLI list TLV	This document
TBD6	Shared BLI TLV	This document

5.2. New Object

IANA is requested to make the assignment from the "PCEP Object" sub-registry as follows:

Value	Name	Reference
TBD3	Traffic Model Object	This document
TBD4	BLI Object	This document

6. Security Considerations

TBD

7. Acknowledgements

8. Normative References

[I-D.geng-spring-sr-enhanced-detnet] Geng, X., Li, Z., and T. Zhou, "Segment Routing for Enhanced DetNet", Work in Progress, Internet-Draft, draft-geng-spring-sr-enhanced-detnet-01, 13 March 2023, <<https://datatracker.ietf.org/doc/html/draft-geng-spring-sr-enhanced-detnet-01>>.

[I-D.ietf-pce-segment-routing-policy-cp] Koldychev, M., Sivabalan, S., Barth, C., Peng, S., and H. Bidgoli, "PCEP extension to support Segment Routing Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-11, 20 June 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-11>>.

[I-D.yzz-detnet-enhanced-data-plane] Geng, X., Zhou, T., Zhang, L., and Z. Du, "DetNet Enhanced Data Plane", Work in Progress, Internet-Draft, draft-yzz-detnet-enhanced-data-plane-02, 24 December 2022, <<https://datatracker.ietf.org/doc/html/draft-yzz-detnet-enhanced-data-plane-02>>.

[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>>.

[RFC4657] Ash, J., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.

[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>>.

- [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>>.
- [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>>.
- [RFC9016] Varga, B., Farkas, J., Cummings, R., Jiang, Y., and D. Fedyk, "Flow and Service Information Model for Deterministic Networking (DetNet)", RFC 9016, DOI 10.17487/RFC9016, March 2021, <<https://www.rfc-editor.org/info/rfc9016>>.

Authors' Addresses

Li Zhang
Huawei

Email: zhangli344@huawei.com

Xuesong Geng
Huawei

Email: gengxuesong@huawei.com

Tianran Zhou
Huawei

Email: zhoutianran@huawei.com