

Network
Internet-Draft
Intended status: Standards Track
Expires: 10 August 2026

Shaofu. Peng
ZTE Corporation
6 February 2026

DetNet EDP Interoperation
draft-peng-detnet-edp-interop-00

Abstract

This document analyzes the interoperation between different EDP solutions. It also suggests which metadata should be used as common fields.

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 10 August 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. Metadata of EDP Solutions	3
3. Interoperation of EDP Solutions	4
3.1. IPv6	5
3.2. MPLS	6
4. IANA Considerations	8
5. Security Considerations	8
6. Acknowledgements	8
7. Normative References	8
Author's Address	9

1. Introduction

[I-D.ietf-detnet-dataplane-taxonomy] defines the categories for classifying the sets of various DetNet EDP (Enhanced Data Plane) solutions. The suitable categories for achieving the goal of deterministic networking are also defined. Similar to the IEEE TSN queueing mechanisms, DetNet also defines several candidate mechanisms based on rate and time respectively. TSN ATS/CBS and DetNet C-SCORE/N-SCORE/gLBF are rate based mechanisms, while TSN TAS/CQF and DetNet EDF/TQF/TCQF/ONTIME are time based mechanisms.

It is known that all TSN queueing mechanisms use the common frame fields, mainly the priority field, or combine it with other classic Ethernet frame fields (such as Destination MAC address, Source MAC address, VID), to map to specific queues. Common fields can bring convenience to the interoperation between different TSN queueing mechanisms. However, the cost is maintaining Per-stream Filtering and Policing (PSFP) on each node. Specifically, the received priority, such as 3 for AVB class A service or 2 for AVB class B service, is mapped to a specific stream gate instance, where one or more IPVs (Internal Priority Value) are further set, and the frame will access the IPV queue according to the rules defined by that stream gate instance. Each IPV queue explicitly configures the queueing mechanism it uses. Different TSN queueing mechanisms will define different styles of PSFPs. In short, interoperation involves two parts of data: simple common fields and complex local states. During interoperation, different queueing mechanisms only need to understand common fields, without needing to know local states related to other mechanisms.

On the contrary, all DetNet EDP solutions choose to carry the state in the packet rather than maintaining the state on the node. Each queueing mechanism has its own particular metadata, which means that when crossing multiple technical domains, multiple metadata may need

to be encapsulated in the packet or switched between multiple metadata at boundary nodes. This seems more complex than the interoperation between different TSN queueing mechanisms. However, the overall system complexity of interoperation for TSN and DetNet is the same. The interoperation still involves two parts of data: simple common fields and complex differentiated fields. During interoperation, different queueing mechanisms only need to understand common fields, without needing to know differentiated fields related to other mechanisms.

This document analyzes the interoperation between different EDP solutions. It also suggests which metadata should be used as common fields. Note that for interoperability, common fields should be stored separately from differentiated fields in the packet header.

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. Metadata of EDP Solutions

[RFC3246] defines a type of PHB called EF (Expedited Forwarding). The suggested DSCP codepoint is 101110. The purpose of EF PHB is to provide building blocks for low packet loss, low latency, and low jitter services. [RFC3246] mentions that the specific details of how to build such services are not within the scope of it. DetNet EDP precisely supplements these details.

If following TSN approach, specific DSCP codepoint (such as 101110) can be mapped to target traffic class with related queueing mechanism configured. Note that a single traffic class may correspond to a set of queues. This will submit the packet to the corresponding queueing mechanism. It requires consistent configuration of all nodes in the domain, which is not a problem when a technical domain is controlled by the same management. This approach seems to find a common fields among all EDP solutions, however, some EDP solutions also have their own particular metadata as additional scheduling parameters. This means that specific metadata already implies a specific queueing mechanism, and even the queueing mechanism type can be explicitly set in the metadata.

The metadata involved may be the following forms:

- * A time slot stack, related period instance, and latency deviation (accumulated by ingress arrival slot deviation and egress departure slot deviation).
- * A single cycle-id, related cycle instance, and optional cycle deviation (accumulated by ingress arrival cycle deviation and egress departure cycle deviation).
- * A single planned residence time, and latency deviation (accumulated by the residence time deviation of each hop).
- * A worst-case latency stack, and latency deviation (accumulated by the queueing delay deviation of each hop).
- * A finish time of previous node, maximum delay of current node, and optional latency deviation that equals to the finish time minus the departure time of previous node.

All EDP solutions essentially adjust the delay of packets at each hop within their technical domain. Therefore, when two technical domains interact, one domain focuses on the latency deviation results of the other domain, rather than the details of adjusting delay within each domain.

The latency deviation should be designed as a common field shared by all EDP solutions, while the rest fields should be treated as differentiated fields specific to each solution. The entry node of a technical domain should perform a damping operation on the received packet based on the latency deviation field in the packet.

3. Interoperation of EDP Solutions

Figure 1 illustrates the scenario of a DetNet path passing through three different technical domains. The controller splits the latency and jitter requirements of the application flow into three parts, and calculates appropriate paths for each technical domain separately. The complete end-to-end path is composed of concatenated paths from three domains, including inter domain links. Assuming that on the unidirectional inter domain link from the upstream domain to the downstream domain, the queueing mechanism of the upstream domain continues to be used.

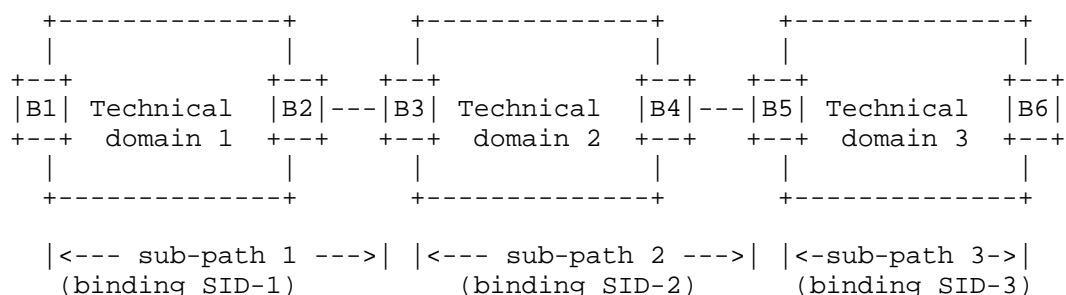


Figure 1: Interoperation Between Technical Domains

At ingress node B1, the application flow is mapped to an end-to-end DetNet path which is represented as <SID-1, SID-2, SID-3> to egress node B6. SID-1, SID-2, and SID-3 bind flow to sub-paths within domain 1, 2, 3, respectively. The FIB states of SID-1, SID-2, and SID-3 also provide necessary information for constructing metadata.

3.1. IPv6

This section describes the interoperation between multiple technical domains by IPv6 encoding [RFC8200].

The packet sent by ingress node B1 may have the following format, where, in the outer IPv6 header, a common field, latency deviation, is set based on the difference between delay bound and actual delay at B1. If the queueing mechanism within domain 1 eliminate jitter per hop, then the latency deviation is contained in Hop-by-Hop Options header. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the latency deviation is contained in Destination Options header.

```

IPv6 header
  [HBH: latency deviation] ;if EDP control jitter per hop
  RH: sub-path 1 with differentiated metadata
  [DOH: latency deviation] ;if EDP only control jitter at exit
IPv6 header
  RH: SID-1, SID-2, SID-3 ; SL = 2
Upper-layer header

```

The latency deviation in HBH may be updated per hop. The latency deviation may also be updated on the last hop. When the packet arrived at the endpoint B3 of sub-path 1, it will be dampened based on the latency deviation, and the outer IPv6 header is removed. Then, the packet continues to be forwarded based on the SID-2 in the next IPv6 header.

The packet sent by node B3 may have the following format, where a common field latency deviation is set based on the difference between delay bound and actual delay at B3. If the queueing mechanism within domain 2 eliminate jitter per hop, then the latency deviation is contained in Hop-by-Hop Options header. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the latency deviation is contained in Destination Options header.

```
IPv6 header
  [HBH: latency deviation] ;if EDP control jitter per hop
  RH: sub-path 2 with differentiated metadata
  [DOH: latency deviation] ;if EDP only control jitter at exit
IPv6 header
  RH: SID-1, SID-2, SID-3 ; SL = 1
Upper-layer header
```

When the packet arrived at the endpoint B5 of sub-path 2, it will be dampened based on the latency deviation, and the outer IPv6 header is removed. Then, the packet continues to be forwarded based on the SID-3 in the next IPv6 header.

The packet sent by node B5 may have the following format, where a common field latency deviation is set based on the difference between delay bound and actual delay at B5. If the queueing mechanism within domain 3 eliminate jitter per hop, then the latency deviation is contained in Hop-by-Hop Options header. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the latency deviation is contained in Destination Options header.

```
IPv6 header
  [HBH: latency deviation] ;if EDP control jitter per hop
  RH: sub-path 3 with differentiated metadata
  [DOH: latency deviation] ;if EDP only control jitter at exit
IPv6 header
  RH: SID-1, SID-2, SID-3 ; SL = 0
Upper-layer header
```

When the packet arrived at the endpoint B6 of sub-path 3, it will be dampened based on the latency deviation, and the outer IPv6 header is removed. The next IPv6 header will also be removed. The packet continues to be forwarded based on the upper-layer header.

3.2. MPLS

This section describes the interoperation between multiple technical domains by MPLS MNA encoding [I-D.ietf-mpls-mna-hdr].

The packet sent by ingress node B1 may have the following format, where, a common field latency deviation is set based on the difference between delay bound and actual delay at B1. If the queueing mechanism within domain 1 eliminate jitter per hop, then the scope of MNA NAS is HBH. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the scope of MNA NAS is Select.

```
... ..
Adj-SID-B2->B3
Node-SID-B3
<NAS: latency deviation, scope=HBH if EDP control jitter per hop
    or scope=Select if EDP only control jitter at exit>
SID-2
SID-3
Upper-layer header
```

The latency deviation in MNA NAS may be updated per hop. The latency deviation may also be updated on the last hop. When the packet arrived at the endpoint B3 of sub-path 1, it will be dampened based on the latency deviation, and the top NAS is removed. Then, the packet continues to be forwarded based on the next label SID-2.

The packet sent by node B3 may have the following format, where a common field latency deviation is set based on the difference between delay bound and actual delay at B3. If the queueing mechanism within domain 2 eliminate jitter per hop, then the scope of MNA NAS is HBH. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the scope of MNA NAS is Select.

```
... ..
Adj-SID-B4->B5
Node-SID-B5
<NAS: latency deviation, scope=HBH if EDP control jitter per hop
    or scope=Select if EDP only control jitter at exit>
SID-3
Upper-layer header
```

When the packet arrived at the endpoint B5 of sub-path 2, it will be dampened based on the latency deviation, and the top NAS is removed. Then, the packet continues to be forwarded based on the next label SID-3.

The packet sent by node B5 may have the following format, where a common field latency deviation is set based on the difference between delay bound and actual delay at B5. If the queueing mechanism within

domain 3 eliminate jitter per hop, then the scope of MNA NAS is HBH. Otherwise, if the queueing mechanism only eliminate jitter at exit of the domain, then the scope of MNA NAS is Select.

```
... ..
Node-SID-B6
<NAS: latency deviation, scope=HBH if EDP control jitter per hop
    or scope=Select if EDP only control jitter at exit>
Upper-layer header
```

When the packet arrived at the endpoint B6 of sub-path 3, it will be dampened based on the latency deviation, and the top NAS is removed. The packet continues to be forwarded based on the upper-layer header.

4. IANA Considerations

There is no IANA requestion for this document.

5. Security Considerations

TBD

6. Acknowledgements

TBD

7. Normative References

[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](https://datatracker.ietf.org/doc/html/draft-ietf-detnet-dataplane-taxonomy-05)>.

[I-D.ietf-mppls-mna-hdr]

Rajamanickam, J., Gandhi, R., Zigler, R., Song, H., and K.
Kompella, "MPLS Network Action (MNA) Sub-Stack
Specification including In-Stack Network Actions and
Data", Work in Progress, Internet-Draft, draft-ietf-mppls-
mna-hdr-19, 3 February 2026,
<[https://datatracker.ietf.org/doc/html/draft-ietf-mppls-
mna-hdr-19](https://datatracker.ietf.org/doc/html/draft-ietf-mppls-mna-hdr-19)>.

- [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>>.
- [RFC3246] Davie, B., Charny, A., Bennet, J.C.R., Benson, K., Le Boudec, J.Y., Courtney, W., Davari, S., Firoiu, V., and D. Stiliadis, "An Expedited Forwarding PHB (Per-Hop Behavior)", RFC 3246, DOI 10.17487/RFC3246, March 2002, <<https://www.rfc-editor.org/info/rfc3246>>.
- [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>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

Author's Address

Shaofu Peng
ZTE Corporation
China
Email: peng.shaofu@zte.com.cn