OPSAWG
Internet-Draft
Intended status: Standards Track
Expires: 29 June 2026

X. Song
Y. Liu
ZTE Corp.

26 December 2025

Export of L4S ECN in IP Flow Information Export (IPFIX)
draft-song-opsawg-ipfix-ecn-00

Abstract

This document defines a set of IP Flow Information Export (IPFIX) Information Elements for monitoring the Low Latency, Low Loss, and Scalable throughput (L4S) service. Specially, these elements enable network operators to monitor the Explicit Congestion Notification (ECN) information of L4S deployment and performance of traffic.

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 29 June 2026.

Table of Contents

1. Introduction

The Low Latency, Low Loss, and Scalable throughput (L4S) service, defined in [RFC9331], introduces a new network service that enables low latency and high throughput for traffic using Scalable congestion controls. To deploy and operate L4S effectively, network operators need visibility into L4S traffic patterns, performance metrics, and interoperability with existing traffic.

IP Flow Information Export (IPFIX) [RFC7011] provides a standard protocol for exporting flow information from network devices. This document defines a set of IPFIX Information Elements specifically designed for monitoring L4S ECN traffic.

These Information Elements are particularly useful during the experimental phase of L4S deployment as specified in [RFC9331], allowing operators to gather data to examine performance and identify nodes where remediation may be necessary to provide the best performance.

2. Terminology

2.1. Terms Used in This Document

This document makes use of the terms defined in [RFC9331], [RFC9330] and [RFC7011].

IPFIX: IP Flow Information Export IPFIX Information Elements

Template: Template Record Options Template: Options Template Record

Data Record Scalable Congestion Control

Classic Service Low Latency, Low Loss, and Scalable throughput (L4S) service

ECN: Explicit Congestion Notification ECT: ECN-capable Transport

Not-ECT: Not ECN-capable transport CE: Congestion Experienced

2.2. 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.

3. Overview of ECN Format

3.1. IPv4 ECN Field

For IPv4 packets, the ECN field is located in the Type of Service (TOS) byte of the IP header, specifically in bits 6 to 7. The ECN field in IPv4 [RFC3168] is showed as follows.

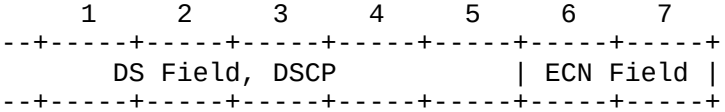

Figure 1: ECN Fields in IPv4

ECN Codepoint values:

00: Not ECT 01: ECT(0)

10: ECT(1) 11: CE

3.2. IPv6 ECN Field

For IPv6 packets, the ECN field is located in the Traffic Class octet, also in bits 6 to 7, as specified in [RFC2474] and [RFC3168].

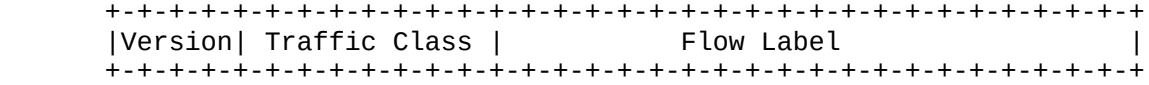

Figure 2: ECN Fields in IPv6

4. IPFIX Information Elements for L4S ECN Monitoring

This section defines the Information Elements for L4S ECN. These elements are intended for experimental use in L4S monitoring.

4.1. ipv4HeaderEcn

Name: ipv4HeaderEcn

ElementID: TBD1

Description: This element is used for capturing the complete ECN state of each packet, enabling detailed analysis of congestion notification. The ECN field is encoded in bits 6 to 7 of the IPv4 TOS byte as defined in [RFC3168]. L4S traffic is identified by the ECT(1) as specified in [RFC9331].

The Information Element encodes only these 2 bits. Therefore, its value may range from 0 to 3.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Additional Information: Refer to the "IPv4 TOS Byte " registry in section 5 of [RFC3168].

Reference: [RFC3168], [RFC9331], this document.

4.2. ipv6HeaderEcn

Name: ipv6HeaderEcn

ElementID: TBD2

Description: This element is used for capturing the complete ECN state of each packet, enabling detailed analysis of congestion notification. The ECN field is encoded in bits 6 to 7 of the IPv6 Traffic Class octet as defined in [RFC3168]. L4S traffic is identified by the ECT(1) codepoint as specified in [RFC9331].

The Information Element encodes only these 2 bits. Therefore, its value may range from 0 to 3.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Additional Information: Refer to the "IPv6 Traffic Class Octet" registry in section 5 of [RFC3168].

Reference: [RFC3168], [RFC9331], this document.

4.3. mplsHeaderEcn

Name: mplsHeaderEcn

ElementID: TBD3

Description: The EXP field of the MPLS label header is used for carring ECN information in the MPLS domain. As recommended in [RFC5129], explicit congestion notification in MPLS should use codepoints in the EXP field.

The Information Element encodes only these 3 bits. Therefore, its value may range from 0 to 7.

It is noted that the information extraction of this information element is only used when the MPLS domain has ECN support.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Additional Information: see [RFC5129] for detailed information for

MPLS ECN tunnel negotiation.

Reference: [RFC5129], this document.

4.4. ipsecSaEcnMode

Name: ipsecSaEcnMode

ElementID: TBD4

Description: The information element indicates whether ECN functionality is allowed for an IPsec Security Association (SA) in tunnel encapsulation mode. The IPsec SA Attribute value 10 is defined for ECN tunnel negotiation as defined in section 9.2.1 of [RFC3168]. The negotiation value includes allowed (value set 1) and forbidden (value set 2) attribute. The allowed value enables ECN congestion notifications and the forbidden value disables such notifications.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Additional Information: See [RFC3168] for detailed information for IPsec tunnel ECN negotiation.

Reference: [RFC5129], this document.

4.5. l2tpEcnNego

Name: l2tpEcnNego

ElementID: TBD5

Description: For L2TP tunnels, ECN processing is performed at the L2TP encapsulation layer. [RFC9601] defines an ECN Capability AVP (Type 103) for negotiation between L2TP Control Connection Endpoints. The presence of this AVP indicates support for ECN propagation.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Additional Information: See [RFC9601] for detailed information for L2TP tunnel ECN negotiation.

Reference: [RFC5129], this document.

4.6. notEctPacketDeltaCount

Name: notEctPacketDeltaCount

ElementID: TBD6

Description: The number of packets since the previous report (if any) in this Flow with ECN codepoint set to Not-ECT (binary 00).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Additional Information: Refer to [RFC9331].

Reference: [RFC3168], [RFC9331], this document.

4.7. notEctPacketTotalCount

Name: notEctPacketTotalCount

ElementID: TBD7

Description: The total number of packets of this Flow with ECN codepoint set to Not-ECT at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Additional Information: Refer to [RFC9331].

Reference: [RFC3168], [RFC9331], this document.

4.8. ect0PacketDeltaCount

Name: ect0PacketDeltaCount

ElementID: TBD8

Description: The number of packets since the previous report (if any) in this Flow with ECN codepoint set to ECT(0) (binary 01).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Additional Information: Refer to [RFC3168].

Reference: [RFC3168], [RFC9331], this document.

4.9. ect0PacketTotalCount

Name: ect0PacketTotalCount

ElementID: TBD9

Description: The total number of packets of this Flow with ECN codepoint set to ECT(0) at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Additional Information: Refer to [RFC3168].

Reference: [RFC3168], [RFC9331], this document.

4.10. ect1PacketDeltaCount

Name: ect1PacketDeltaCount

ElementID: TBD10

Description: The number of packets since the previous report (if any) in this Flow with ECN codepoint set to ECT(1) (binary 10).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Additional Information: Refer to [RFC3168].

Reference: [RFC3168], [RFC9331], this document.

4.11. ect1PacketTotalCount

Name: ect1PacketTotalCount

ElementID: TBD11

Description: The total number of packets of this Flow with ECN codepoint set to ECT(1) at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Additional Information: Refer to [RFC9331].

Reference: [RFC3168], [RFC9331], this document.

4.12. cePacketDeltaCount

Name: cePacketDeltaCount

ElementID: TBD12

Description: The number of packets since the previous report (if any) in this Flow with ECN codepoint set to CE (Congestion Experienced, binary 11).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Additional Information: Refer to [RFC9331].

Reference: [RFC3168], [RFC9331], this document.

4.13. cePacketTotalCount

Name: cePacketTotalCount

ElementID: TBD13

Description: The total number of packets of this Flow with ECN codepoint set to CE at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Additional Information: Refer to [RFC9331].

Reference: [RFC3168], [RFC9331], this document.

4.14. l4sCeMarkRatioDelta

Name: l4sCeMarkRatioDelta

ElementID: TBD14

Description: The proportion of L4S packets marked with the CE codepoint, calculated over the reporting interval since the previous report (if any). This element represents the incremental CE marking rate for L4S traffic within the Flow. It is calculated as:

CE-marked L4S packets (delta count) / Total L4S packets (delta count)

Where L4S packets are those identified by the ECN codepoint ECT(1). The result is a ratio ranging from 0.0 to 1.0.

Abstract Data Type: float32

Data Type Semantics: quantity

Additional Information: This ratio provides a near-real-time view of congestion dynamics for L4S traffic. It is useful for detecting transient congestion events and monitoring short-term performance.

Reference: [RFC3168], [RFC9331], this document.

4.15. l4sCeMarkRatioTotal

Name: l4sCeMarkRatioTotal

ElementID: TBD15

Description: The proportion of L4S packets marked with the CE codepoint, calculated over the total lifetime of the Flow. This element represents the total CE marking rate for L4S traffic within the Flow. It is calculated as:

CE-marked L4S packets (total count) / Total L4S packets (total count)

Where L4S packets are those identified by the ECN codepoint ECT(1). The result is a ratio ranging from 0.0 to 1.0.

Abstract Data Type: float32

Data Type Semantics: quantity

Additional Information: This ratio provides a long-term view of

congestion experienced by L4S traffic. A value of 0.0 indicates no congestion marking, while a value approaching 1.0 indicates persistent or severe congestion.

Reference: [RFC9331], this document.

5. Operational Considerations

For IPsec tunnels, monitoring ECN requires exporting both outer and inner IP header ECN fields (ipHeaderOuterEcn and ipHeaderInnerEcn), along with ipsecSaEcnMode (see section 4.4). Relying solely on the outer IP header ECN field may be insufficient, as it could be set to Not ECT due to tunnel mode restrictions. Similarly, for L2TP tunnels, ECN monitoring should be verified the l2tpEcnNego element (see section 4.5) except the ECN information extraction from tunnel outer header and inner header of packets.

For MPLS tunnels, the ECN handling mechanism differs fundamentally from IP based tunnels. ECN information is not carried in a dedicated IP header field but is encoded within the MPLS label stack using the EXP field, as defined in [RFC5129]. Therefore, monitoring ECN over MPLS requires exporting the mplsHeaderEcn element defined in Section 4.3. This element captures the congestion indication as conveyed within the MPLS domain, which is independent of the inner IP packet’s ECN field.

When measuring the proportion of packets marked with the CE codepoint, the CE marking rate for L4S traffic should be calculated specifically for flows identified as ECT(1) (L4S traffic identifier) prior to marking, rather than aggregating all CE-marked packets irrespective of their original ECT codepoint. This ensures the performance of L4S services can be accurately monitored and distinguished from Classic ECN traffic, which may have different congestion response characteristics.

6. Security Considerations

The security considerations for IPFIX [RFC7011] apply to this document. The elements for ECN reveal information about endpoint ECN capabilities. Although the information may generally be not sensitive, operators should consider applicable privacy regulations. IPFIX records containing L4S monitoring information SHOULD be transported using secure protocols such as TLS or DTLS and satisfy the mutual authentication between IPFIX Exporting Processes and IPFIX Collecting Processes as specified in [RFC7011].

7. IANA Considerations

IANA is requested to assign the following Information Elements in the IPFIX Information Elements registry.

Element ID	Name	Reference
TBD1	IPv4HeaderEcn	This document
TBD2	IPv6HeaderEcn	This document
TBD3	MPLSHeaderEcn	This document
TBD4	ipsecSaEcnMode	This document
TBD5	l2tpEcnNego	This document
TBD6	notEctPacketDeltaCount	This document
TBD7	notEctPacketTotalCount	This document
TBD8	ect0PacketDeltaCount	This document
TBD9	ect0PacketTotalCount	This document
TBD10	ect1PacketDeltaCount	This document
TBD11	ect1PacketTotalCount	This document
TBD12	cePacketDeltaCount	This document
TBD13	cePacketTotalCount	This document
TBD14	l4sCeMarkRatioDelta	This document
TBD15	cePacketTotalCount	This document

Table 1: New IPFIX Information Elements

8. References

8.1. Normative References

[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>.

[RFC3168] Ramakrishnan, K., Floyd, S., and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, DOI 10.17487/RFC3168, September 2001, <https://www.rfc-editor.org/info/rfc3168>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <https://www.rfc-editor.org/info/rfc7011>.

[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

[RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC 2474, DOI 10.17487/RFC2474, December 1998, <https://www.rfc-editor.org/info/rfc2474>.

[RFC5129] Davie, B., Briscoe, B., and J. Tay, "Explicit Congestion Marking in MPLS", RFC 5129, DOI 10.17487/RFC5129, January 2008, <https://www.rfc-editor.org/info/rfc5129>.

[RFC9330] Briscoe, B., Ed., De Schepper, K., Bagnulo, M., and G. White, "Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture", RFC 9330, DOI 10.17487/RFC9330, January 2023, <https://www.rfc-editor.org/info/rfc9330>.

[RFC9331] De Schepper, K. and B. Briscoe, Ed., "The Explicit Congestion Notification (ECN) Protocol for Low Latency, Low Loss, and Scalable Throughput (L4S)", RFC 9331, DOI 10.17487/RFC9331, January 2023, <https://www.rfc-editor.org/info/rfc9331>.

[RFC9601] Briscoe, B., "Propagating Explicit Congestion Notification across IP Tunnel Headers Separated by a Shim", RFC 9601, DOI 10.17487/RFC9601, August 2024, <https://www.rfc-editor.org/info/rfc9601>.

Authors' Addresses

Xueyan Song
ZTE Corp.
Email: song.xueyan2@zte.com.cn

Yao Liu
ZTE Corp.
Email: liu.yao71@zte.com.cn