

Export of ECN information in IPFIX

[draft-song-opsawg-ipfix-ecn-01](#)

Xueyan Song **ZTE Corp.**

Yao Liu ZTE Corp.

Overview

- **Low latency, Low Loss and Scalable Throughput (L4S)** is an end-to-end technology that is designed to enable both high throughput and low latency even under network congestion. The L4S compatible end-to-end congestion control supports fine-grained feedback from ECN marking.
- The purpose of this document is to define a set of **IP Flow Information Export (IPFIX)** Information Elements for L4S ECN monitoring, enabling network administrators/operators with access to IP Flow information to examine network performance and identify network nodes where congestion is experienced and network optimization is necessary.

Positioning

- **Features**

- Defines a set of IPFIX IEs for L4S ECN monitoring.
- Provides operational considerations for L4S ECN information export.

- **Benefits**

- Provides network visibility and diagnostics, and quantifies the network congestion status.
- Monitors the consistency of ECN configuration in IPsec/L2TP tunnels, benefiting diagnostics of ECN propagation issues.
- Provides IPFIX IEs for L4S ECN marking export and packets statistics to facilitate adjustments of network traffic patterns and supports network planning.

- **Values**

- Uniform export information helps multi-vendor interoperability.
- IPFIX monitoring improves the operation and management of L4S service.

- **Gap coverage**


- Meets the L4S deployment and fills the gap for L4S ECN monitoring using IPFIX.

Monitoring Logic


- In **Metering Process**, uses **ECN Field as Flow Key** and defines an **L4S ECN flow** as a set of packets that have common traffic properties (such as source IP address, destination IP address) and have the ECN field set to ECT(1).
 - The ECN field is located in the IP header and is a typical example of a packet header field. According to Section 2 of RFC7011, packet header fields can serve as Flow Keys to define a flow, and can also be included as measured properties in a Flow Record.
- In **Exporting Process**, encapsulates a set of IPFIX Information Elements, including ECN-related fields state and statistical data for ECN information export of traffic flows.
 - The Exporter sends IPFIX messages containing flow records to the Collector.
- By analyzing the observed data, the **Collector** evaluates which Observation Points in the network support L4S (i.e., can correctly propagate ECN marks), and assesses whether the end-to-end ECN marking behavior meets expectations (e.g., whether congestion points mark packets with CE).





ECN Information Export – data plane



RFC 3168, 9331 ECN field in IP header



RFC 5129, 3270 EXP field in MPLS header



Information Elements defined:

- ipv4HeaderEcn
- ipv6HeaderEcn
- notEctPacketDeltaCount
- notEctPacketTotalCount
- ect0PacketDeltaCount
- ect0PacketTotalCount
- ect1PacketDeltaCount
- ect1PacketTotalCount
- cePacketDeltaCount
- cePacketTotalCount
- **mplsHeaderEcn**

draft-halmer-mpls-ecn, ECN Field in ISD/PSD/MNA

Feedback and Updates

- **1 - Questions for clarification and correction**
- Gregory Mirsky (Target of IPFIX notification and observation location):
 - **Clarification:** The observation point for ECN information can be flexibly positioned depending on the monitoring purpose.
 - **Update Made:**
 - Basic ECN fields (IPv4/IPv6/MPLS headers) can be extracted at any IPFIX-capable node along the path.
 - Tunnel ECN negotiation status IEs are only available at the tunnel endpoints.
 - Cumulative statistics IEs may be prioritized at aggregation or egress nodes.
- Gorry Fairhurst (Reference correction):
 - **Issue:** Noted that RFC 2474 supersedes older definitions of the IPv4 TOS octet and IPv6 Traffic Class octet.
 - **Update Added:** Corrected references throughout (Sections 3.1, 3.2, 4.1, 4.2) to align with the DiffServ Field as defined in RFC 2474.

Feedback and Updates

- **2 - Questions for performance and scalability**
- Joel Halpern (L4S identification and export)
 - **Issue Raised:**
 - How to distinguish L4S traffic from classic ECN/non-ECN traffic?
 - How to avoid the overhead of exporting ECN information for every packet across many flows?
 - **Updates Made (this version):**
 - Addressed overhead concern: Added text recommending the use of **sampling methods (RFC5475)** to preserve statistical representativeness while reducing export overhead.
 - **Planned Updates (future version):**
 - Will address traffic distinction: Plan to add filtering guidance.
 - Proposed method: Use packet **filtering (RFC5475)** at the Observation Point to select packets with the ECT(1) codepoint.

Feedback and Updates

- **3 – Questions for IE discussion**
- Sebastian Moeller & Ingemar Johansson S
 - **Issue 1: Probability of CE Marking**
 - Feedback: Noted that whether CE marking is probabilistic or deterministic is a function of the AQM implementation, not the marking law (e.g., DCTCP vs. RFC 3168).
 - Update: Added clarification that in L4S deployments, CE marking may be either probabilistic or deterministic, depending on the AQM. The defined IEs for statistical counts of CE-marked packets enable quantitative monitoring and management, typically based on flow data acquired via sampling.
 - **Issue 2: Ambiguity of Ratio IEs**
 - Feedback: Pointed out that "I4sCeMarkRatioTotal" and "I4sCeMarkRatioDelta" are useful but could introduce ambiguity. Proposed to calculate these ratios from base metrics instead.
 - Update: Removed the two ratio IEs from the draft. Added text clarifying that calculating derived metrics (e.g., CE marking ratios) from base counters (e.g., ect1PacketTotalCount, cePacketTotalCount) is an implementation issue for the Collector and is out of scope for this draft.

Next Steps

- Further develop the IEs details and elaborate on implementation considerations.
- Continue to update the draft based on feedback.