

OPSAWG
Internet-Draft
Intended status: Standards Track
Expires: 16 June 2026

Jinjie. Yan
ZTE Corporation
13 December 2025

Export of Energy Consumption Information in IPFIX
draft-yan-opsawg-ipfix-energy-consumption-00

Abstract

This document introduces new IPFIX IEs for exporting energy consumption information of physical entities in a network device. New Information Elements are defined to report instantaneous and average energy consumption information at device, line-card, and port granularity.

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 16 June 2026.

Copyright Notice

Copyright (c) 2025 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. Terminology	3
3. New IPFIX IEs for Energy Consumption Information	3
3.1. deviceRealtimePower	3
3.2. deviceAveragePower	4
3.3. lineCardRealtimePower	4
3.4. lineCardAveragePower	4
3.5. portRealtimePower	5
3.6. portAveragePower	5
4. Use Cases	6
4.1. Energy-Aware Routing	6
4.2. Per-Flow Energy Consumption Monitoring	6
5. Operational Considerations	7
6. Security Considerations	8
7. IANA Considerations	8
8. Acknowledgements	8
9. References	8
9.1. Normative References	8
Author's Address	9

1. Introduction

Energy consumption has emerged as a critical operational metric in modern data centers and communication networks, driven by increasing demands for sustainability and cost efficiency. Operators require accurate, traffic-correlated power information to support applications such as energy-aware routing, per-service energy cost accounting, and fine-grained power management.

Traditional monitoring protocols, like SNMP, rely on periodic polling, which often fails to capture the dynamic interplay between traffic loads and power draw, leading to delayed or inaccurate insights. In contrast, the IP Flow Information Export (IPFIX) protocol [RFC7011] offers a proven, event-driven framework that inherently binds metering observations to traffic events at the source, ensuring timely and causal reporting.

To address these needs, this document defines six new Information Elements (IEs) that enable an IPFIX Exporting Process to report energy consumption for physical entities within a network device, including the entire device, line cards, and ports. These IEs support both instantaneous (realtime) power values, which provide snapshots at the moment of export for immediate analysis, and average power values over the measurement interval since the last export, which are useful for trend detection and aggregated accounting. The

measurements are exported via existing IPFIX triggering mechanisms—such as packet count thresholds, active timeouts, and inactive timeouts—preserving causality between observed traffic and reported energy metrics.

This document introduces new IPFIX IEs for exporting energy consumption information of physical entities in a network device to facilitate energy-aware networking by providing granular, real-time data for optimized network operations.

1.1. 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 [RFC2119].

2. Terminology

This document makes use of the terms defined in [RFC7011]:

- * IPFIX
- * IPFIX Information Elements
- * Exporting Process
- * Metering Process
- * Template Record
- * Data Record
- * Exporter
- * Collector

3. New IPFIX IEs for Energy Consumption Information

The following new Information Elements are defined:

3.1. deviceRealtimePower

Name : deviceRealtimePower

ElementID: TBD1

Description: Instantaneous total power consumption of the entire exporting device at the time of export.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: watts

Reference: This document.

3.2. deviceAveragePower

Name : deviceAveragePower

ElementID: TBD2

Description: Average total power consumption of the entire exporting device over the interval since the last export of this Template.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: watts

Reference: This document.

3.3. lineCardRealtimePower

Name : lineCardRealtimePower

ElementID: TBD3

Description: Instantaneous power consumption of the line card at the time of export.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: milliwatts

Reference: This document.

3.4. lineCardAveragePower

Name : lineCardAveragePower

ElementID: TBD4

Description: Average power consumption of the line card identified over the interval since the last export.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: milliwatts

Reference: This document.

3.5. portRealtimePower

Name : portRealtimePower

ElementID: TBD5

Description: Instantaneous power consumption of the port at the time of export.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: milliwatts

Reference: This document.

3.6. portAveragePower

Name : portAveragePower

ElementID: TBD6

Description: Average power consumption of the port over the interval since the last export.

Abstract Data Type: unsigned32

Data Type Semantics: quantity

Units: milliwatts

Reference: This document.

4. Use Cases

4.1. Energy-Aware Routing

In this scenario, network controllers may require accurate, traffic-correlated power consumption metrics of physical links (represented by ports) to compute paths that minimize total energy cost. Traditional periodic polling mechanisms often fail to capture the dynamic relationship between instantaneous traffic load and actual power draw, leading to suboptimal routing decisions.

By leveraging the IPFIX-based energy telemetry defined in this document, a network device can export multi-level power consumption data (e.g., port, line card, and device levels) triggered by traffic events such as packet count thresholds or active/inactive timeouts. The exported data records can provide both instantaneous and average power values over the observation interval, enabling the controller to derive meaningful energy cost metrics, such as per-link energy-per-bit or average power under load, for use in path computation algorithms. This tight coupling between traffic behavior and energy reporting ensures that routing decisions reflect real-time energy efficiency characteristics of network links.

4.2. Per-Flow Energy Consumption Monitoring

For fine-grained energy accounting and auditing of specific high-bandwidth ("elephant") flows, such as backup traffic between data center servers, it is essential to attribute the energy consumed across multiple physical components (e.g., ingress/egress ports and their respective line cards) to the flow itself.

Using the mechanisms described in this document, an IPFIX Exporting Process can associate energy consumption measurements with a specific IP flow (identified by, e.g., a 5-tuple) and trigger synchronized reporting when predefined traffic thresholds are met. By including physical entity identifiers (e.g., lineCardId, portId) and the newly defined energy Information Elements, the exporter can generate multiple correlated data records corresponding to the various physical interfaces involved in forwarding the flow. This enables collectors to reconstruct a complete, multi-path energy footprint of the flow across the device, supporting accurate per-flow energy cost attribution and operational analysis.

5. Operational Considerations

The export and interpretation of energy consumption information in IPFIX are intended to be driven by traffic activity observed at specific points in the data plane. Implementations SHOULD configure monitoring policies that associate one or more physical entities—such as a device, line card, or port—with an observation object (e.g., an interface or a flow).

Energy telemetry records SHOULD be generated when predefined traffic-driven conditions are met. These conditions MAY include:

- * The number of packets observed since the last export exceeds a configured threshold;
- * An active timeout elapses under continuous but low-volume traffic; or
- * An inactive timeout elapses after no packets are observed for a configured duration.

The Exporting Process collects power measurements—such as instantaneous and average power—for each associated physical entity. These measurements are encoded into Data Records using the newly defined energy Information Elements (e.g., lineCardRealtimePower, portAveragePower). Each Data Record SHOULD also include identifier Information Elements (e.g., lineCardId, portId) to unambiguously associate the reported energy values with their corresponding hardware components.

The resulting Data Records are exported to a Collector within an IPFIX Message, along with a reference to the applicable Template Record (e.g., via Template ID).

A Collector receiving such messages SHOULD possess the corresponding Template Record to correctly parse the structure and semantics of the Data Records.

The Collector parses each Data Record to recover:

- * The identifier(s) of the physical entity (or entities) linked to the observation object; and
- * The associated energy consumption information.

This end-to-end mechanism ensures that energy reports are causally tied to traffic activity.

6. Security Considerations

TBA

7. IANA Considerations

This document requests IANA to create a new IE under the "IPFIX Information Elements" registry [RFC7012] available at [IANA-IPFIX].

Element ID	Name	Reference
TBD1	deviceRealtimePower	Section 3.1
TBD2	deviceAveragePower	Section 3.2
TBD3	lineCardRealtimePower	Section 3.3
TBD4	lineCardAveragePower	Section 3.4
TBD5	portRealtimePower	Section 3.5
TBD6	portAveragePower	Section 3.6

Table 1: IPFIX Information Elements Registry

8. Acknowledgements

TBA

9. References

9.1. Normative References

[IANA-IPFIX]

IANA, "IP Flow Information Export (IPFIX) Entities",
<<https://www.iana.org/assignments/ipfix>>.

[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>>.

- [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>>.
- [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model
for IP Flow Information Export (IPFIX)", RFC 7012,
DOI 10.17487/RFC7012, September 2013,
<<https://www.rfc-editor.org/info/rfc7012>>.

Author's Address

Jinjie Yan
ZTE Corporation
China
Email: yan.jinjie@zte.com.cn