

FROM PRINCIPLES TO STANDARDS

A JOURNEY TOWARD ENERGY-AWARE STREAMING

GreenIO London 2026 · Marisol Palmero · September 2026

<https://standards.ieee.org/>
<https://europe.ieee.org/>

**"The standard takes an *hour* to write
and five years to agree on."**

The hour is the easy part.

Five years in working groups taught me this.

Getting people with competing interests to agree on the same words — that is the hard part. That is the work.

WHO DECIDES WHAT "GOOD" LOOKS LIKE?



The Journey Toward Energy-Aware Streaming


From an Industry Connection Activity to create an IEEE Standard Proposal(PAR) and why that matters

() Project Authorization Request (IEEE SA)*

AI compute overtakes video streaming around 2025

TWh per year · 2000–2030 · Dashed = projected

— Video Streaming - - AI / ML Compute ● Crossover ~2025



Sources: IEA Electricity 2024; Goldman Sachs Power Up (2024); Malmodin & Lundén (2018); Cisco VNI; Epoch AI; Lawrence Berkeley National Lab. AI figures = training + inference workloads in data centres. Video streaming = end-to-end (network + DC). 2026–2030 projections are Goldman Sachs estimates. All figures are illustrative approximations.

Energy Location: Video Streaming Vs. AI Inference - *Structural inversion*

Energy location: video streaming vs AI inference



Video Streaming

60–75% of total system energy is at the *consumer's device* — the TV, phone, or laptop. The data centre (CDN, encoding, servers) accounts for only 25–40%.

Implication: the primary carbon lever is **display technology and default settings**. Codec efficiency and server optimisation move a minority of the total.



AI Inference

97–99% of total inference energy is in the *data centre* — GPU/TPU clusters and their cooling overhead. The user's device (typing a query, reading the response) accounts for <0.1%.

Implication: the primary carbon lever is **data centre hardware efficiency and grid carbon intensity**. User behaviour and device type are largely irrelevant.

The governance gap is parallel — neither domain has per-session carbon attribution. But the solutions are completely different: streaming needs device policy, AI needs data centre policy.

- > Telling users “to watch less Netflix” affects the dominant energy layer.
- > Telling users to “ask fewer AI questions” has negligible effect because the energy is entirely in the data centre regardless.

THREE CONVERGING PRESSURES

Why energy measurement in “streaming” can no longer be voluntary or estimated

REGULATORY DIVERGENCE

**CSRD, SEC Climate
Disclosure Rule, UK SDR, ...**

US retreats from voluntary ESG

Europe mandates full supply chain

reporting — Omnibus I in force

Terminology is no longer enough.

Evidence is.

DATA QUALITY IS PUBLIC

>70%

of data is estimated, not measured

Omnibus I splits primary/secondary

data for the first time

>70% of reported sustainability

data estimated via proxies

AI analysis only works on real data.

NO SHARED GOVERNANCE

20+

frameworks; none interoperable

No standard governs exchange

of sustainability data across

value chain boundaries

Result: duplicated effort,

unverifiable claims, blind spots.

WHAT IS AN IEEE SA INDUSTRY CONNECTIONS ACTIVITY?

© Not yet a standard

A pre-competitive collaboration space

Open to any organisation — no voting rights required

Builds the consensus that a PAR needs to succeed

Think of it as: institutional R&D before the standard

⬡ Why it matters

Creates the evidence base regulators later cite

Attracts early adopters before the standard is final

De-risks the PAR ballot — consensus built in advance

IEEE SA credential without 5-year standard timeline

The IC is where you prove the problem is real and the community is ready.

The PAR is what you submit when you have that proof.

IC Activity

Open collaboration



PAR Approval

IEEE SA ballot



Working Group

Drafts standard



Ballot

Member vote



Published

IEEE Standard

B2B TRUSTWORTHY SUSTAINABLE VALUE CHAIN

The IEEE Industry Connections initiative — and why this PAR is part of something larger

B2B TSVC — Industry Connections Activity

TRACK 1 Standard

Trustworthy Energy Attribution & Measurement (TEAM)

BTSC - Primary sponsor
Green ICT - Secondary sponsor

Next PAR deadline: 19 Oct 2026


TRACK 2 Implementation

Sustainable & Trustworthy
AI Governance
Certification

TRACK 3 Adoption

Implementation Partners
Cross-sector rollout
Interop & compliance

ENERGY DATA ACROSS THE STREAMING VALUE CHAIN LACKS A COMMON TRUST FRAMEWORK



No standard exists to establish the interoperability, provenance, integrity, or verifiability of energy data exchanged between any of these operators.

The problem is not just data quality. It is the absence of trust infrastructure.



**"Standards are how we make
individual / organization insights
into collective action."**

Marisol Palmero · m.palmeroamador@ieee.org

LEAVING YOU WITH THIS THOUGHT...

HOW STANDARDS PROLIFERATE:
(SEE: A/C CHARGERS, CHARACTER ENCODINGS, INSTANT MESSAGING, ETC.)





What GoS brings to a Normative Contribution

One methodology, offered as technical input and validation data — not the mandated method

THE OWL CORE · Online Watt Lab, published methodology, live

1 Measured at the wall, not modelled

Real hardware, real workloads, external smart-plug sampling at 1 Hz. Idle baseline captured before every run; result = $\Delta W \times \Delta t$.

2 Every run self-calibrates

Per-run baseline, cooldown gate, focus mode, drift term. Noise floor characterised per rig, not assumed.

3 Every result carries a confidence flag

- Repeatable $\geq 95\%$, ≥ 10 polls
- Early insight $\geq 80\%$, ≥ 5 polls
- Need more data — published, not hidden.

4 Evidence travels with the number

Raw per-second and baseline samples, thermal trace, hardware fingerprint, tool and method version, in one JSON envelope.

5 Two lenses, declared

Marginal (idle subtracted) and attributional (idle shared) reported side by side, so the boundary is explicit.

Three years · encode, delivery and playback measured on the same plug family · OWL bench and REM field readers agree within 0.2 %

CONTRIBUTES



as WG technical
input

WHAT WE CAN CONTRIBUTE TO THE STANDARD

QUALIFIED — traffic lights

→ PAR 5.2 c, d · Layer 2

- ✓ A confidence class derived from measured noise, not asserted: tier = $f(\text{signal} / \text{uncertainty, sample count})$.
- ✓ An uncertainty recipe any segment operator can run: calibrated noise + per-run variance + drift.
- ✓ Minimum evidence to retain so a green flag can be re-checked by a third party.

COMPARABLE — apples to apples

→ PAR 5.2 a, b, e · Layers 3–4

- ✓ A boundary declaration: what is inside the meter, which lens (marginal / attributional), how idle is handled.
- ✓ Minimum context fields per record: hardware fingerprint, tool + methodology version, duration, poll count.
- ✓ A running JSON result envelope (v1, live since Aug 2026) as a seed for the EMR; a cross-instrument agreement test as a template for verification status.

METHODOLOGY-NEUTRAL BY DESIGN: OWL is one input among several. The Working Group decides what becomes normative.