APRIL 28, 2026

The EED reporting obligation: what does it really deliver?

Since 2024, the European Energy Efficiency Directive (EED) has required data centers above 500 kW to report annually on energy consumption, water use, PUE and residual heat utilization. The first results are disappointing: only 36% of European data centers reported, data quality is variable and crucial metrics such as IT utilization, GPU idle consumption and water consumption per AI run are completely missing. Enforcement is actually non-existent.

At the same time, the rise of liquid cooling — direct-to-chip and immersion cooling — undermines the already limited value of PUE as an efficiencymetric. Liquid cooling lowers the PUE, but at the same time changes the numerator andthe denominator of the calculation, making comparison with air-cooled environments virtually impossible. While the industry enthusiastically communicates PUE gains, absolute energy consumption is increasing due to higher power densities.

The conclusion is sharp: the EED measures what is measurable, not what matters. As long as IT utilization, GPU efficiency, and system output are outside the reporting

framework, sustainability reporting remains an infrastructure audit — not evidence of actual efficiencyimprovements.

- What is needed: Extension of the EED framework to include the Server Idle Coefficient (SIC), mandatory
- GPU reporting, a TUE equivalent for liquid-cooled environments and real-world enforcement.

The EED reporting obligation: the framework on paper

Since 2024, European data centers with an installed IT capacity of at least 500 kW have been required to report annually to a European database. The obligation stems from Article 12 of the revised Energy Efficiency Directive (EED, 2023/1791) and the accompanying Delegated Regulation 2024/1364. On paper, it is ambitious: energy consumption, water consumption, PUE, temperature settings, waste heat utilization, use of renewable energy and — from 2025 — also IT equipment capacity for newly installed servers.

The intention is clear. The European Commission wants to gain insight into the environmental impact of data centers, wants to be able to make targeted policy and wants to be able to set minimum performance standards if the data give cause to do so. The first reporting period ran until 15 September 2024 (for calendar year 2023). From 2025, May 15 will be the annual deadline.

It sounds like a solid foundation. The reality is less encouraging.

The first results: a sober balance sheet

The European Commission commissioned an evaluation report based on the data reported in 2024. The outcome is telling: only 36% of the data centers in the EU that are subject to the obligation actually submitted a report. That amounts to 770 locations of an estimated more than 2,100 in scope.

In the Netherlands, the picture was no better. Of the 160 data centers that had to report to the RVO (Netherlands Enterprise Agency), 27 left crucial fields empty — almost exclusively parties with American ownership. The RVO does not have legal instruments to enforce additional information. The result: empty forms without consequences.

In its evaluation, the European Commission acknowledges that the quality and completeness of the reported data 'could be improved' and makes proposals to make the system 'clearer, more reliable and easier to use'. These are friendly descriptions for a system that is structurally inadequate in its first year.

Wat de EED verplicht te meten	Wat structureel ontbreekt
Totaal energieverbruik (kWh)	IT-bezettingsgraad / serverbenutting
PUE (Power Usage Effectiveness)	Server Idle Coefficient (SIC)
Waterverbruik (WUE)	Waterverbruik per AI-inferentie of trainingsrun
Restwarmtebenutting (ja/nee)	Hoeveelheid werkelijk teruggeleverde restwarmte (MWh)
Gebruik hernieuwbare energie (%)	Directe vs. certificatengebaseerde groene stroom
IT-apparatuurcapaciteit (nieuw, ≥ 2025)	GPU-idle-verbruik en thermische throttling
Temperatuurinstellingen	Energie per eenheid IT-output (werk/joule)

The table shows the pattern: the EED measures what is visible in terms of infrastructure. What happens at the IT level — utilization, waste, GPU behavior — remains out of the picture. This is no coincidence. It is a reflection of how the sector has presented itself for years: as an infrastructure provider, not as a manager of digital output.

PUE in the EED: A trusted metric with known shortcomings

PUE is the central metric in the EED reporting framework. That's understandable: it's the best-known, most standardized, and widely available indicator of data center efficiency. But anyone who has read my previous analysis of PUE knows that this choice poses a fundamental problem.

PUE measures the ratio of total data center consumption to IT consumption. A PUE of 1.2 means that 20% of the total energy is spent on supporting systems — cooling, UPS, lighting. That is infrastructural efficiency. It says nothing about what the IT equipment actually does with that energy.

A data center with 35% server utilization and a PUE of 1.15 scores better in the EED framework than a data center with 80% utilization and a PUE of 1.35. That is untenable from a system point of view. The first environment structurally wastes IT capacity; the second uses its infrastructure efficiently. PUE does not make that distinction — and neither does the EED.

In my previous articles, I introduced the Server Idle Coefficient (SIC) as a necessary addition: a metric that measures what proportion of the energy consumption of IT equipment takes place while no functional output is delivered. Without a SIC — or an equivalent — the EED remains structurally blind to the biggest inefficienciesin modern data centers: underutilization and waste.

Liquid cooling: efficiency gains or metric pollution?

While the EED puts PUE at the center, a technological development that is currently being embraced at a rapid pace further undermines the usefulness of that metric: liquid cooling.

Direct-to-chip cooling, immersion cooling and rear-door heat exchangers are no longer a thing of the future. They have become mainstream, driven by the explosive growth of GPU clusters for AI workloads. Racks with capacities of 80 to 120 kW per rack — compared to 5 to 15 kW in traditional enterprise IT — are no longer manageable with air cooling. The market for liquid cooling in AI data centers is growing at a rate of more than 30% per year.

The cooling gain is real. Research published by NVIDIA and Vertiv shows that liquid cooling can reduce total data center consumption by approximately 10%, while improving Total Usage Effectiveness (TUE) by more than 15%. Immersion cooling achieves PUE values of 1.02 to 1.05 — impressive on paper.

The PUE problem in liquid cooling

But this creates a metrics problem that the industry largely ignores. Liquid cooling not only changes the denominator of the PUE calculation (the IT consumption), but also the numerator (the total consumption), in a way that makes comparison with air-cooled environments impossible.

With conventional air cooling, the boundary between 'IT consumption' and 'facility consumption' is clear: the server consumes power, the CRAC unit consumes power, and the PUE is the ratio. With liquid cooling, these limits are intertwined. A heat exchanger that is directly connected to the chip — is that IT consumption or facility consumption? Pumping and circulation capacity for coolant are considered IT burden in some methodologies, and overhead in others. This makes PUE comparison between liquid- and air-cooled data centers actually apples and oranges.

The scientific research into the HAWK supercomputer (published in ScienceDirect, 2025) shows this concretely: a lower PUE does not have to go hand in hand with lower total energy consumption in liquid cooling. The optimal setting depends on the outside temperature, the supply temperature of the cooling water and the mix between liquid and air cooling in the system. In other words, a lower PUE may be the result of a parameter choice, not a real efficiency improvement.

Illustratie: hetzelfde datacenter, twee PUE-uitkomsten

Stel: een datacenter met een IT-last van 5 MW en vloeistofkoeling.
Scenario A: koelpomp- en circulatievermogen geclassificeerd als 'facilitair' → PUE 1,08
Scenario B: koelpomp- en circulatievermogen geclassificeerd als 'IT-last' → PUE 1,02

Het absolute energieverbruik is in beide scenario's identiek.
De gerapporteerde PUE verschilt 6 procentpunt — louter op basis van definitiekeuze.

De EED schrijft geen eenduidige classificatiemethode voor vloeistofkoeling voor.
Dat betekent dat gerapporteerde PUE-waarden voor vloeistofgekoelde omgevingen onderling en ten opzichte van luchtgekoelde omgevingen niet vergelijkbaar zijn.

Higher density, higher absolute consumption

There is another fundamental point that receives little attention in the sector. Liquid cooling enables higher power densities — and that opportunity is being exploited. Racks that used to run at 10 kW now run at 80 to 120 kW. The result: absolute IT consumption per rack increases dramatically, while PUE decreases. The total energy consumption of the data center increases in net terms — but the PUE improvement presented gives the impression of efficiencygains.
This is not hypothetical reasoning. It's at the heart of the AI energy paradox: better cooling technology enables higher density, higher density requires more energy, and the metric that is supposed to measure the result — PUE — becomes less reliable because of that same technology.

What the EED should add now

In its evaluation report (October 2025), the European Commission proposed to improve the reporting system, including a focus on minimum performance standards. That is the right direction. But the proposals get stuck in the familiar framework: more of the same, better measurement, more uniform measurement.
What is missing is a fundamental expansion of the metric framework. Concrete steps are needed on three points:

IT utilisation as a mandatory indicator

Energy consumption without a link to IT output is meaningless. The EED should require data centres to report on their average server utilisation – the percentage ofinstalled IT capacity that is actually doing active work. The Server Idle Coefficient (SIC) provides an operationally applicable framework for this.

A data center that runs 40% of its servers structurally idle while the energy meter continues to run is not efficient — regardless of the PUE. As long as the EED does not measure that, half the story is missing.

GPU reporting as a separate category

GPU environments behave fundamentally differently than CPU-driven data centers. They have extreme peak power, strong load fluctuations and significant idle consumption — even outside of training runs. The current EED obligation treats all IT equipment as onecategory. This leads to structural bias in environments where GPU clusters are dominant.

Mandatory GPU reporting — including idle consumption, thermal throttling, and energy per inference — is not a luxury. It is a prerequisite for policy that is in line with the reality of AI infrastructure.

A TUE equivalent as an additional metric for liquid cooling

Total Usage Effectiveness (TUE) — the ratio of total data center energy to energy used by actual compute units (CPU, GPU, memory) — provides a more robust alternative to PUE in liquid-cooled environments. This makes pump and circulation losses visible and makes apples-to-apples comparison between cooling technologiespossible.

The EED should mandate TUE as an additional mandatory metric for data centers above a certain power density, or for environments where liquid cooling handles more than 30% of the IT burden.

Enforcement: the Achilles heel of the system

The EED does not include mandatory penalties for non-reporting data centres. Member States are allowed to set up their own enforcement mechanisms — but practice shows that in most countries this has not been worked out, if at all. In the Netherlands, the RVO has no legal instrument to enforce incomplete reports. The situation is similar in other Member States.

The result: of the data centers that are required to report, 64% do not report. Of those who do report, a significant proportion provide incomplete data. Large American hyperscalers leave fields empty — and are not affected by it at all. Parliamentary questions in the House of Representatives about the behavior of Microsoft and Google only resulted in a ministerial answer four months later.

A reporting obligation without enforcement is a voluntary disclosure. That is fundamentally different from what the EED intended.

What this means for the sector

The conclusion is not that the EED is worthless. The obligation to report on energy consumption, water use and PUE is a step in the right direction. It forces data centers to

think about their metering infrastructure, stimulates uniformity and provides the European Commission with a basis for policy.

But she measures what is measurable, not what matters. As long as IT utilisation, GPU efficiency and system output are outside the scope, EED reporting remains an infrastructure audit – not evidence of real efficiencyimprovements.

For the sector, the following applies: those who wait for minimum obligations and then report exactly what the law requires, miss the opportunity to tell the story themselves. The companies that already voluntarily report on SIC, GPU utilisation and actual returned residual heat are building the legitimacy that the sector needs to obtain permits, maintain support and substantiate grid congestion discussions with facts.

Transparency is not a compliance exercise. It's a strategic choice — and for a sector that calls itself critical infrastructure, it's also a moral imperative.