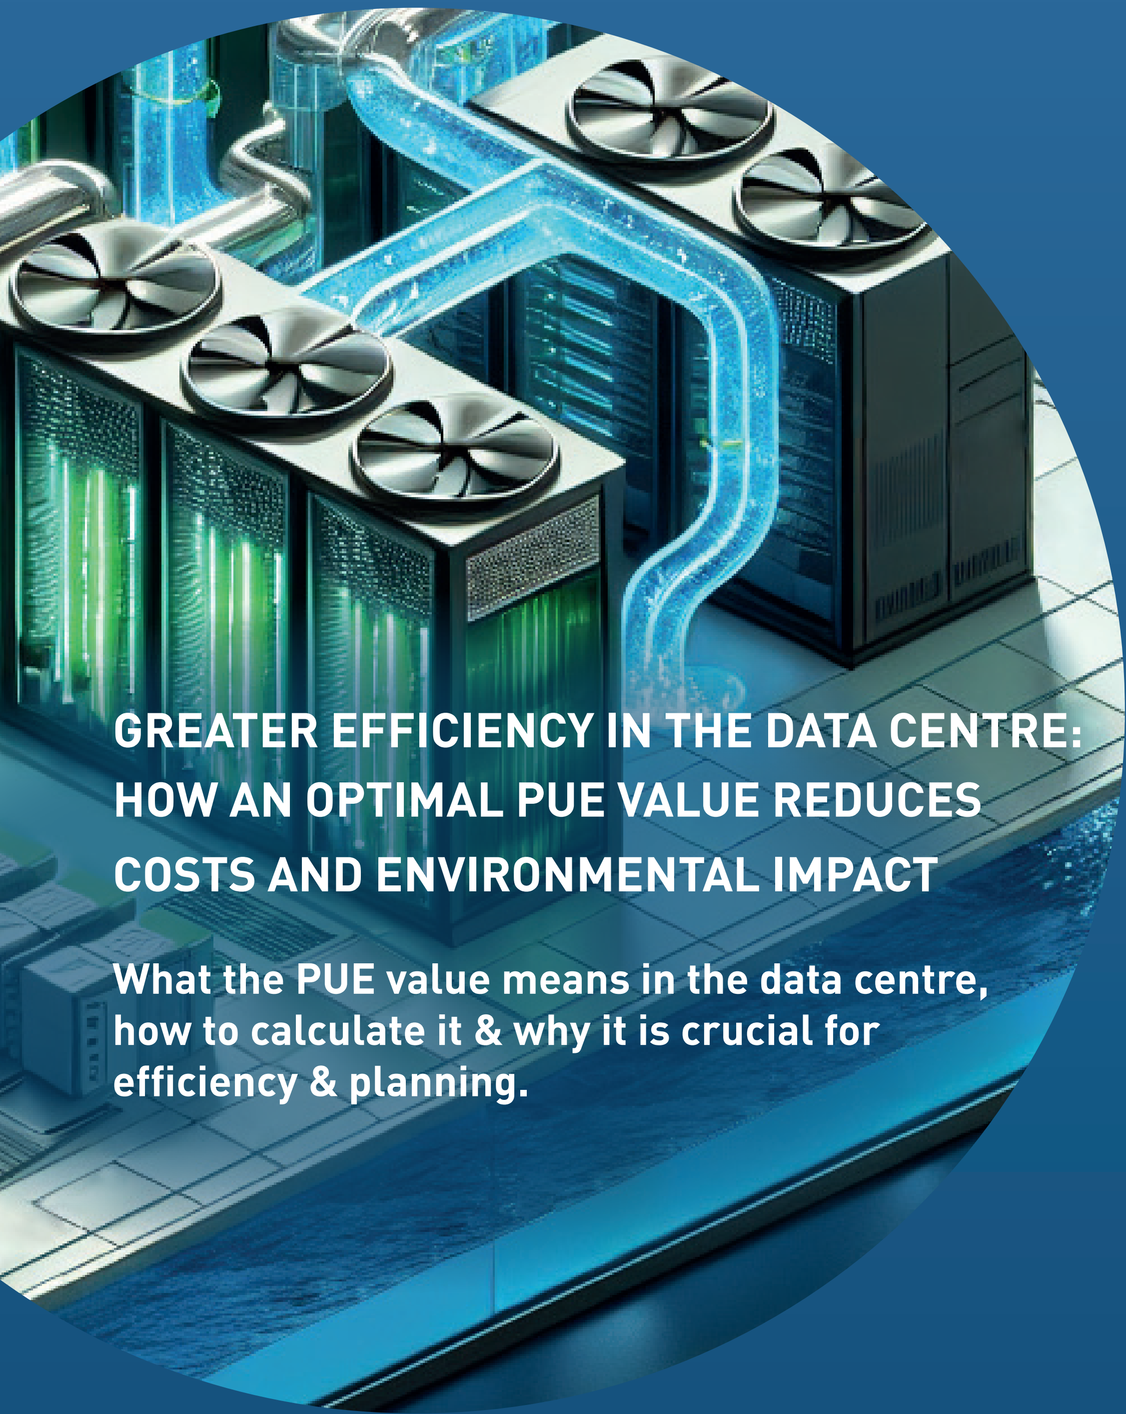

GREATER EFFICIENCY IN THE DATA CENTRE: HOW AN OPTIMAL PUE VALUE REDUCES COSTS AND ENVIRONMENTAL IMPACT

What the PUE value means in the data centre, how to calculate it & why it is crucial for efficiency & planning.

We are part of the solution__

Table of contents

Energy efficiency in data centres: The significance of PUE

An important benchmark for objectively assessing the energy efficiency of data centres is the PUE (Power Usage Effectiveness) index. It shows how the total energy consumption of a data centre - including servers, storage, and network equipment - compares to the energy used by the IT systems. In other words, how much energy goes to IT and how much is needed for cooling, lighting, or power distribution. But why is PUE so central to data centres?

Data centres are now among the fastest-growing electricity consumers worldwide, making them a key lever for the energy transition in the digital age. By 2023, their global electricity demand was already between 300 and 380 terawatt hours (TWh), which corresponds to 1 to 1.3 % of global consumption.

The International Energy Agency (IEA) expects this figure to more than double, reaching around 945 TWh by 2030, driven by the growing use of artificial intelligence and data-intensive applications. The industry faces a dual challenge: reducing energy consumption while meeting rising demand.

With rising electricity demand, energy efficiency requirements in data centres are becoming increasingly stringent. Electricity can account for up to 60 % of a data centre's operating costs (OPEX), while investors and customers are demanding credible ESG strategies. In addition, regulators such as the EU are making efficiency metrics mandatory.

This whitepaper analyses:
- What the PUE value indicates and where it comes from
- How it can be calculated and influenced
- What potential for optimisation exists
- Why PUE remains indispensable

What does the PUE value say about a data centre?

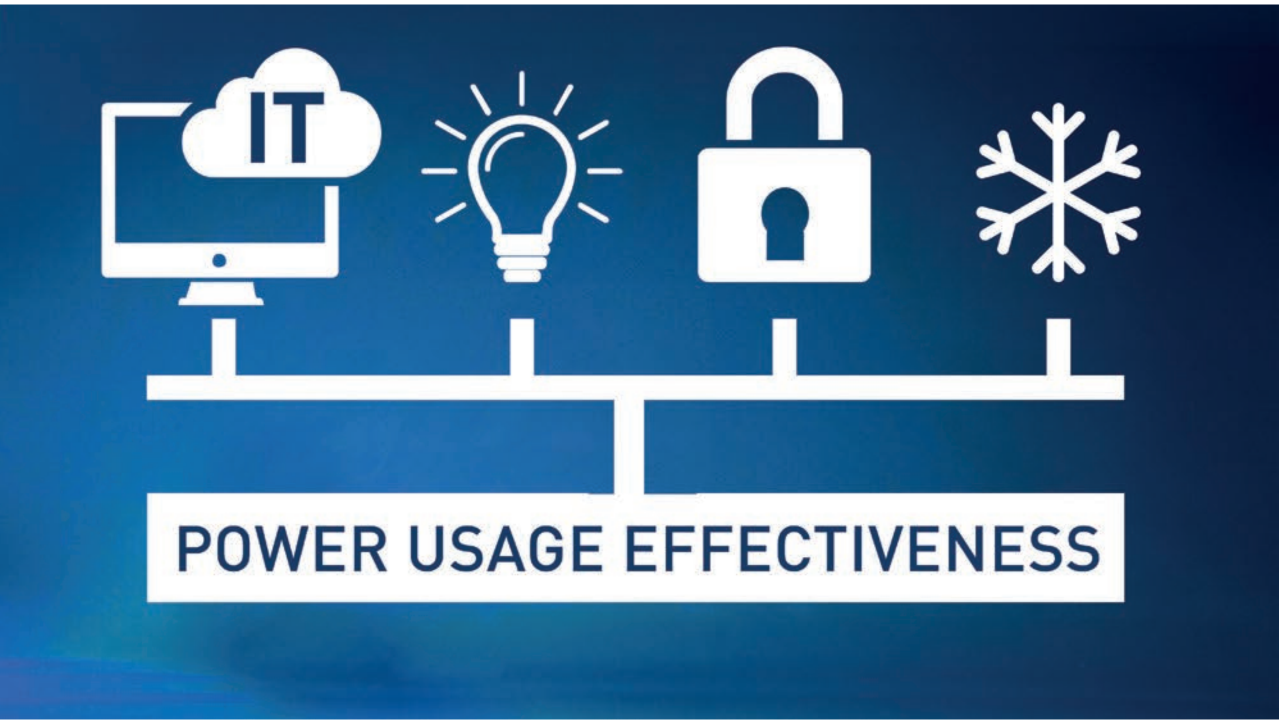

PUE (Power Usage Effectiveness) is the globally recognised standard for comparing the energy efficiency of data centres. It shows how much of the energy consumed is actually used by IT, and how much is consumed by supporting infrastructure such as cooling, uninterruptible power supply (UPS), lighting, or security.

The benefit of PUE for data centres lies in its simplicity:

- Clear guidance for operators, investors, and regulators
- Opportunity to compare locations, technologies, and measures
- Serves as a basis for benchmarks and continuous optimisation

Origin and standardisation

The PUE was introduced in 2007 by The Green Grid consortium with the aim of making energy efficiency uniformly measurable on a global scale. It is now firmly established in international standards:

- **ISO/IEC 30134-2:2016**
  Globally recognised standard
- **EN 50600-4-2**
  European standard for measuring the energy efficiency of data centres

These standards also define measurement methods, measurement points, and quality requirements.

PUE in the context of other metrics

The PUE only reflects the ratio of energy flows in a data centre. In other words, the PUE shows how much of the energy consumed is actually used for IT systems. However, this is also its limitation: PUE provides no information about actual computing power, the climate impact of the electricity used, or the resource consumption of cooling technology in the data centre. To address these gaps, three additional key metrics are commonly used in practice:

DCiE (Data Centre Infrastructure Efficiency)
is the reciprocal of PUE. While PUE quantifies infrastructure overhead, DCiE describes the percentage of total electricity consumption attributable to IT: a positive perspective on the same issue.

WUE (Water Usage Effectiveness)
adds the dimension of water consumption to the PUE of a data centre. Water requirements can be considerable, especially with adiabatic cooling systems or open recoolers. The WUE shows how many litres of water per kilowatt hour are used to produce IT power – an essential factor in water-critical regions.

CUE (Carbon Usage Effectiveness)
measures CO_2 emissions per IT kWh, addressing the biggest limitation of PUE: the energy source. A data centre with a low PUE but high fossil fuel usage may ultimately have a larger carbon footprint than one powered by green electricity, even if its PUE is technically higher.

These key metrics are becoming increasingly important, particularly in the context of ESG strategies, carbon footprints, and regulatory reporting. They transform PUE from an isolated efficiency indicator into a key component of holistic sustainability monitoring.

How can the PUE value be calculated?

Calculating the PUE value follows a simple principle, but the underlying measurement methodology is standards-based and more detailed. The key factor is clearly distinguishing between the power consumption of IT systems and the total consumption of the entire data centre.

The formula for PUE is:

$$\text{PUE} = \frac{\text{Total electricity consumption of the data centre}}{\text{Electricity consumption of IT}}$$

This formula is simple but revealing. An ideal PUE value of 1.0 in a data centre means that all of the energy generated flows directly into the IT systems. Not a single watt is lost to air conditioning or infrastructure. Technically, this ideal state cannot be achieved; it serves as a theoretical benchmark.

Total energy consumption (counters):
Includes all systems required to maintain IT operations, for example:

- Cooling systems (air cooling, liquid cooling, adiabatic systems, etc.)
- Losses in UPS systems and transformers
- Lighting
- Access control, fire protection, and security technology

IT load (denominator):
Includes only systems that directly provide IT services, for example:

- Server
- Storage systems
- Network components
- KVM switches
- Monitors and IT consoles

Sample calculation:

An example data centre consumes a total of 2,000,000 kWh of electricity per year, of which 1,000,000 kWh is required for IT.
PUE = 2,000,000 / 1,000,000 = 2.0
This means that only 50 % of the energy is actually used for IT.
The rest goes towards supporting infrastructure.

Measurement methods
the PUE according to ISO/IEC 30134-2

The significance of a PUE value depends largely on the measurement method used. The international standard ISO/IEC 30134-2 defines four levels for this purpose, depending on where the IT power consumption is measured.

PUE 0

The simplest variant is PUE0. In this case, no direct measurements are taken; the value is estimated using empirical data or rough assumptions. This method is quick but imprecise and is suitable only for initial estimates.

PUE 1

PUE1 is based on measurements taken behind the uninterruptible power supply (UPS). This is the standard in many data centres, as the IT load can be easily measured at this point. However, conversion losses and subordinate consumers are often not precisely accounted for.

PUE 2

PUE2 offers greater accuracy, with measurements taken directly behind the power distribution units (PDUs), at the point where power is delivered to the server racks. This method provides more realistic data but requires more sophisticated instrumentation.

PUE 3

PUE3 achieves the highest measurement quality. Here, data is collected directly from the individual IT components, for example via rack-level metering or socket meters. This variant is considered particularly accurate, but is technically and organisationally complex.

For reliable comparisons, many standards require year-round measurements, known as trailing twelve months (TTM). This compensates for seasonal variations in cooling and energy requirements. Shortterm or project specific evaluations are also possible – for example, using **iPUE** (Interim PUE) for temporary measurements, or **pPUE** (Partial PUE) when only certain sub-areas, such as individual server rooms or racks, are considered.

Typical PUE values in practice

In practice, the PUE value varies considerably depending on the year of construction, technology, location conditions, and utilisation of a data centre.

While modern hyperscalers achieve peak values, older medium-sized data centres sometimes operate significantly above this.

Data centre type	Typical PUE value	Practical application
Hyperscalers (Google, Microsoft)	1.08 - 1.20	Google (TTM-PUE: 1.09); Microsoft design target: 1.12
Modern colocation providers	1.3 - 1.5	e.g. Equinix, Digital Realty, according to the Uptime Institute
Medium-sized data centres	1.6 - 2.0	Often without free cooling, lower utilisation
Legacy sites	> 2.5	In the case of outdated infrastructure or low IT load
Locations in cold regions	≈ 1.1	e.g. in northern regions with free air cooling (airside economiser)

Factors influencing actual PUE values

Cooling strategy: Free cooling in cold regions can significantly improve PUE.

IT utilisation: The better servers and storage are utilised, the lower the relative overhead.

Economies of scale: Large data centres can systematically exploit efficiency potential.

Design standards: New data centres with hot aisle containment, optimised airflow, or liquid cooling achieve better values.

Why is the PUE value of a data centre important?

The PUE value is more than just a technical metric. It is a measure of energy efficiency, a business lever, a sustainability indicator, and increasingly relevant for regulatory purposes. For operators, investors, and customers, it serves as a key reference point in the planning, operation, and evaluation of a data centre.

Economic relevance

Energy efficiency has a direct impact on the operating costs of a data centre. A PUE of 1.6 means that 60 % of the energy is consumed by infrastructure.

If the PUE drops to 1.4, this overhead is reduced by over 12 %. For a 10 MW data centre, this corresponds to around 1.25 GWh of energy saved annually, with corresponding savings in electricity costs, cooling technology, and grid fees.

Sample calculation

IT load: 10 MW
PUE reduction: from 1.6 to 1.4
Savings: 2 MW × 24/7 operation = 1.25 GWh/year

Why PUE is becoming increasingly important beyond technology

A low PUE not only reduces the energy consumption of a data centre but also directly impacts its carbon footprint. In climate strategies such as Science-Based Targets (science-based emission targets) or the EU Taxonomy (EU classification system for sustainable economic activities), PUE has long served as an efficiency indicator that makes progress quantifiable. It also plays a central role in emissions reporting, for example for Scope 2 emissions (indirect emissions from purchased energy). In countries with a high CO_2 factor in their electricity mix, such as Poland or Ireland, every PUE improvement significantly enhances a data centre's climate performance.

Efficiency indicators become legally binding
With the Climate Neutral Data Centre Pact, the EU is increasingly anchoring PUE in regulations. New locations in the Netherlands, for example, may only be approved with a PUE ≤ 1.2.

PUE as a strategic factor in management, ESG, and purchasing
From a stakeholder perspective, PUE is also gaining strategic importance. Operators use it as an operational control parameter and as a KPI (key performance indicator) in their energy management systems. Investors integrate PUE values into ESG due diligence processes (sustainability assessments in the context of investment decisions) and financing models, such as green bonds (sustainability-linked bonds). For customers, it is increasingly becoming part of tendering and certification criteria. Hyperscalers such as Google and Microsoft publish it as an integral part of their sustainability reports – and link their net-zero targets (decarbonisation by 2040/2050) to it.

PUE as part of regulatory requirements minimum standards
Relevant standards and certifications, such as EN 50600-4-2, ISO / IEC 30134-2, LEED (Leadership in Energy and Environmental Design), and the EU Code of Conduct for Data Centres, establish PUE limits as minimum requirements. A low PUE therefore not only improves a data centre's operating cost structure but also enhances its regulatory assessment and market value. Efficiency therefore becomes a licensing factor - ecologically, economically, and reputationally.

Factors influencing the PUE value

The PUE value is the result of a complex interplay of technical, infrastructural, and operational factors.

How cooling technology significantly influences PUE
Cooling is the largest infrastructure energy consumer in a data centre and therefore one of the key factors in reducing PUE. The more efficiently heat is dissipated, the lower the proportion of non-IT consumers in the total energy requirement.

Classic air cooling,
supported by free cooling or adiabatic evaporation systems, is widely used but reaches its physical and energetic limits as power density increases. Direct evaporation remains in use in older systems but is significantly less efficient.

Liquid cooling,
particularly direct-to-chip liquid cooling (DLC), is considered groundbreaking. It dissipates waste heat directly at the processor surface using a circulating coolant. This significantly reduces the energy required for air conditioning while enabling higher packing densities. Forecasts predict that DLC will become the leading cooling technology by 2030, not only because of its thermal efficiency but also because it is a prerequisite for achieving low PUE values in AI and HPC environments. For operators who consistently focus on efficiency, DLC is a central component of future data centre developments.

Other factors influencing PUE

In addition to cooling technology, other technical and structural factors determine the PUE value of a data centre.
In energy distribution, the efficiency of the uninterruptible power supply (UPS) plays a central role. Modern online UPS systems achieve efficiencies of over 97 %. Losses occur primarily in older transformers or in unnecessary conversion stages, which can be largely avoided by using direct current distribution. Other factors:

Building concept:
The building design also has a measurable impact on energy efficiency. Consistent hot/cold aisle containment, short cable lengths, and well-insulated enclosures minimise thermal losses. Modular construction methods also support demand-based scaling while maintaining high space efficiency.

Location:
The location itself is another influencing factor. In cold and dry regions, the outdoor climate significantly reduces cooling requirements, with corresponding positive effects on PUE. At the same time, high electricity prices and CO_2 intensities in the grid reinforce the ecological and economic benefits of any efficiency improvements.

Capacity utilisation:
The utilisation of the IT infrastructure is not directly recorded by the PUE, but it does influence it indirectly. When server utilisation is low, the energy requirements of the infrastructure remain virtually constant, while the useful energy decreases, thereby worsening the PUE. Virtualisation, intelligent load balancing, and AI-based workload management can be used to increase IT utilisation in a targeted manner, thereby improving the PUE without having to replace hardware.

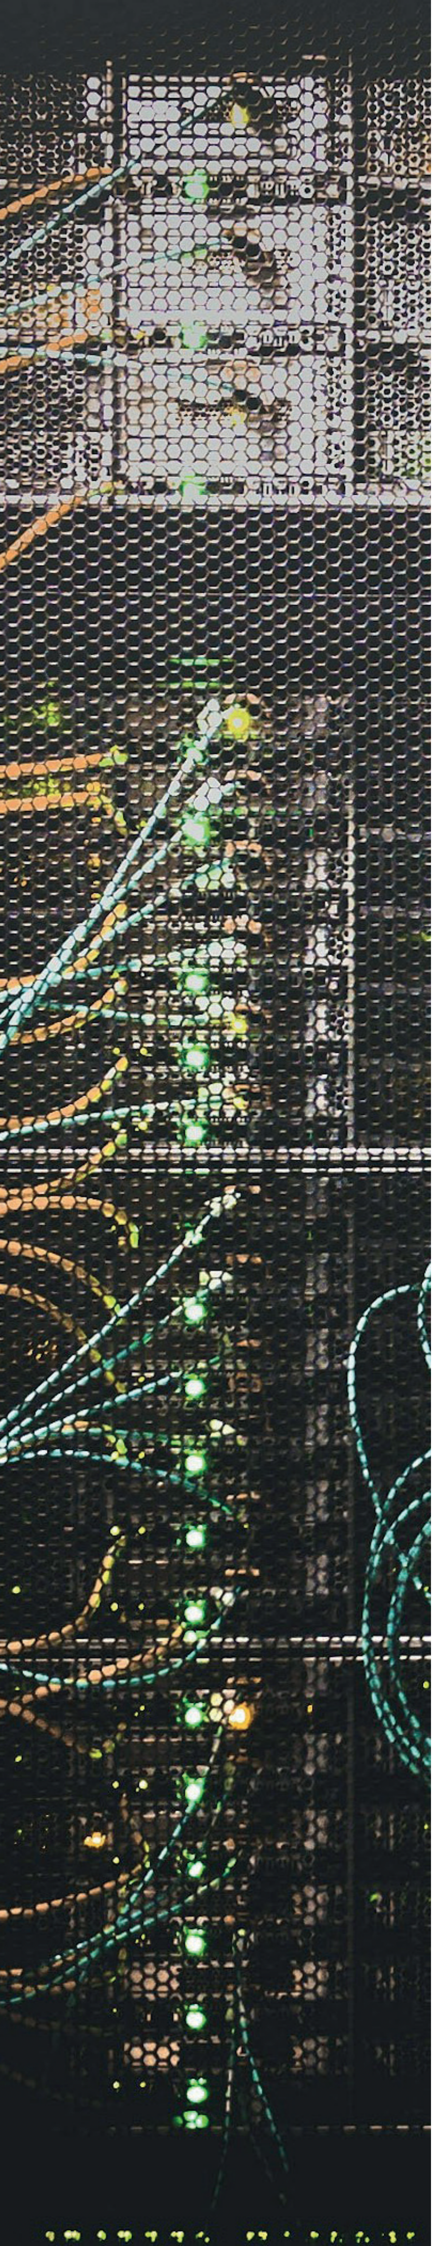

Optimise PUE value:
Measures for more efficient data centres

Low power usage effectiveness (PUE) is not the result of individual measures, but rather the systematic interaction of technical, infrastructural, and organisational decisions. Anyone who wants to set up data centres efficiently should consistently implement the following factors when constructing new buildings or modernising existing ones.

1. Modernising cooling technology
Air conditioning remains the dominant factor influencing PUE. Data centre operators should focus on efficient systems such as hot aisle containment, adiabatic cooling, free outdoor air utilisation, or direct-to-chip liquid cooling (DLC). DLC in particular enables low PUE values with increasing rack density and power-hungry workloads such as AI and HPC.

2. Strategically integrate waste heat
Feeding the waste heat from a data centre into local district heating networks does not improve the PUE, but it does improve the overall energy efficiency of the data centre. The EN 50600 standard reflects this effect via the Energy Reuse Factor (ERF). Operators should consider technical recycling as an integral part of their ESG strategy, especially in urban environments.

3. Consolidate IT infrastructure
Low utilisation leads to disproportionately high energy consumption. Virtualisation, consolidated workloads, and load shifting allow unused resources to be shut down without compromising computing power. This indirectly lowers the PUE and reduces the base load of the supporting systems.

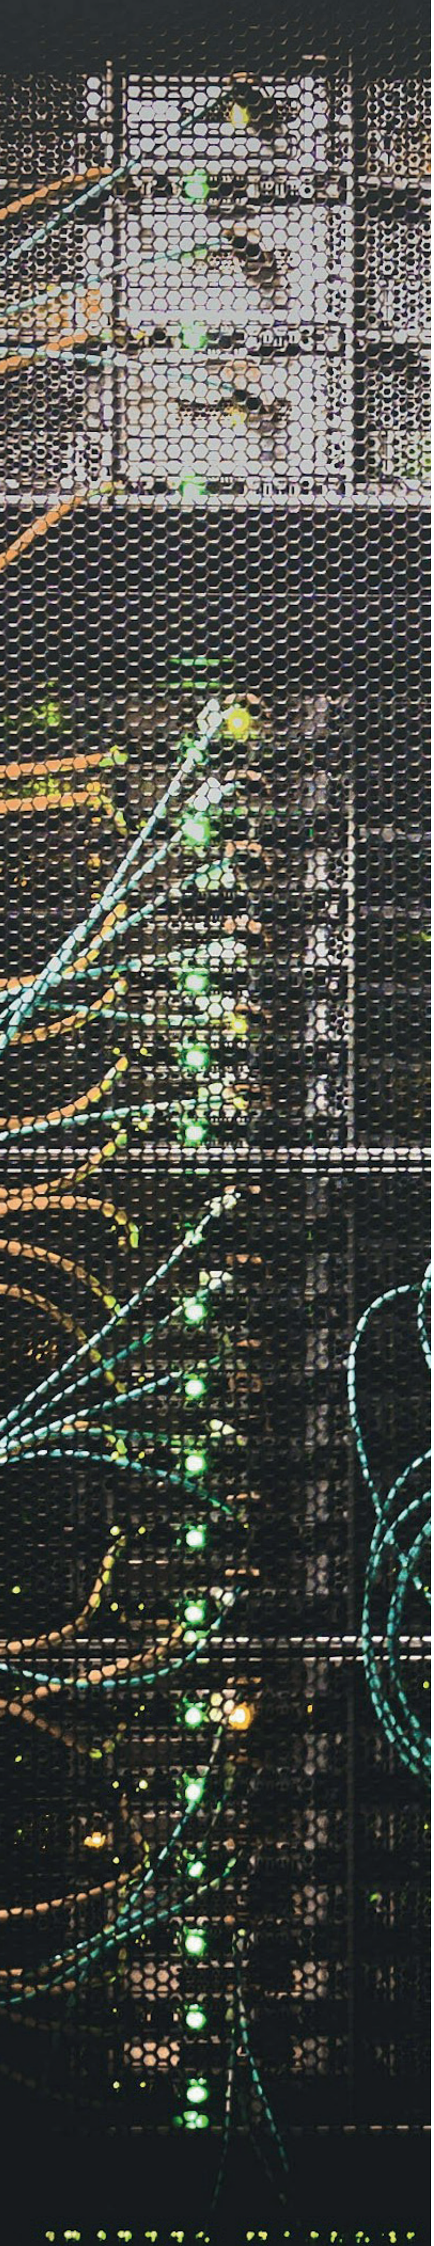

4. Automate monitoring

Energy efficiency cannot be controlled without transparency. Data centres should establish real-time PUE dashboards and AI-supported control systems for fans, UPS, and cooling systems. Studies show that automation reduces infrastructure consumption by up to ten per cent – continuously and at scale.

5. Design robust infrastructure

Liquid-based cooling systems require piping infrastructures that are low-maintenance, corrosion-free, and temperature-resistant. Systems such as aquatherm blue based on PP-RCT, offer a long-term, reliable solution here. The Novva Data Centre (USA) reference shows how high-quality infrastructure creates the technical basis for PUE target values below 1.24.

Where PUE reaches its limits

Power Usage Effectiveness (PUE) is a globally established benchmark for the infra-structural energy efficiency of data centres. It shows how much energy actually benefits IT and how much is used for supporting systems such as cooling or power distribution. However, its significance has clear limitations.

The PUE does not measure the efficiency of the IT systems themselves. It does not account for how much computing power is actually generated per kilowatt hour. Two data centres with identical PUEs can perform very differently, depending on hard-ware generation, utilisation, or degree of virtualisation.

The PUE does not provide any information about the origin of the electricity. Whether fossil fuels or renewable energy are used is not taken into account. A data centre with a very low PUE may have a worse carbon footprint than one with a higher PUE, depending on the energy mix used.

The PUE can be influenced. By selecting measurement points, excluding certain consumers or considering seasonal periods, a value can be obtained that has little to do with actual overall efficiency.

The PUE ignores conflicting objectives. Higher levels of redundancy in data centres – such as Tier III or IV – structurally increase energy requirements.
Greater availability inevitably means lower efficiency. Anyone who looks at the PUE in isolation is ignoring these strategic considerations.

PUE needs context

PUE remains relevant as a technical guideline value - but it needs context. Additional metrics such as Carbon Usage Effectiveness (CUE), Water Usage Effectiveness (WUE), and Energy Reuse Factor (ERF) complete the picture. Only their combination allows a reliable assessment of a data centre's efficiency, sustainability, and future viability. In short, PUE is a starting point - not an end point.

Why cooling represents the greatest opportunity for PUE

Cooling and air movement often account **30 – 45 %** for a significant portion of a data centre's energy consumption. The shift towards direct-to-chip (DLC) liquid cooling for AI workloads (racks with 50 to 120 kW) increases demands on the entire infrastructure - especially on piping systems, which are now a constant source of friction losses, inefficiency, and potential contamination.

The stainless steel standard and its downsides

304 stainless steel pipes form the backbone of most DLC systems, chosen for their low purchase cost, wide availability, high pressure and temperature resistance, and global approval. However, practical and laboratory data suggest a different picture:

Metal leaching:
Independent fluid tests from several US hyperscale locations detected manganese levels up to 25 times higher than the EPA drinking water limits, as well as elevated magnesium and molybdenum levels - all attributable to stainless steel pipes.

Corrosion under real conditions:
Research published in Nature Materials shows that glycol coolants (25 to 50 per cent), with pH values of 6.5 to 8.5 and flow rates of up to 4.8 m^3/h - standard for modern DLC systems - accelerate pitting corrosion in 304 stainless steel and undermine conventional failure predictions.

Result:
Increased ion contamination necessitates more frequent adjustments to water treatment and reduces the service life of coolants from the expected 7–10 years to only 3–5 years in high-performance applications.

The true costs of "cheaper" infrastructure

A typical 3,785-litre DLC cycle calculated over 10 years:

Stainless steel system:

Coolant replacement every 3 – 5 yrs: € 66,000

Filter replacement every 6 - 12 months: € 17,500

Maintenance, testing, repairs: € 52,500

Total: € 136,000

Alternative polymer system:

Coolant replacement every 7 – 10 yrs: € 26,500

Filter replacement every 2 – 3 yrs: € 1,800

Minimum maintenance: € 2,800

Total: € 31,100

Savings per 3,785 litres over 10 years: € 104,900

Cost projection based on data centre megawatts

DLC-cooled sites typically consume 9,460 to 11,350 litres of coolant per MW. This is what the projection looks like:

IT load	Coolant volume	10-year costs stainless steel	10-year costs polymer	Forecast Savings
1 megawatt	9,460 – 11,350 litres	€ 339,000 – 409,000	€ 78,000 – 94,000	€ 262,000 – 315,000
10 megawatts	94,600 – 113,500 litres	€ 3.39 – 4.09 million	€ 0.78 – 0.94 million	€ 2.62 – 3.15 million

For a 10 MW site (≈94,600 litres of coolant), switching to corrosion-free piping can save around € 2.62 million over 10 years – in maintenance and replacement costs alone.

Conclusion

The PUE metric is considered a key benchmark for efficiency in the data centre industry. However, real improvements can only be achieved with an infrastructure that not only delivers short-term investment savings but also supports your operational expenditure (OPEX) and availability goals.
Experience shows that cheap pipes often turn out to be a cost trap - hidden contaminants, higher filter costs, and the need to replace the coolant earlier than expected frequently result in unexpected additional expenses. Operators committed to efficiency should therefore take a close look at their entire cooling circuit - especially at what is inside their pipes.

If you are planning a new data centre or want to operate more efficiently by modernising your existing one, please contact us. Our team of experts at aquatherm will support you with sound advice, planning, and design.

Sources and information

Google Environmental Report 2024
IEA Data Centres Report
Nature Materials Corrosion Research

The information contained in this white paper, in particular the costs, refer to the year 2025.
Further information available on request.

aquatherm®

Made in Germany

TÜV Rheinland ZERTIFIZIERT
Management System
ISO 9001:2015
ISO 14001:2015
ISO 50001:2018
www.tuv.com
ID 0091005348

aquatherm GmbH
Biggen 5 | 57439 Attendorn | Germany
Tel: +49 2722 950 0 | Email: info@aquatherm.de
www.aquatherm.de