Beyond PUE:

A New Class of Data Center Metrics

By Nicola Phillips,
Copywriter at Soluna

October 26, 2022

Table of Contents

Power Usage Effectiveness (PUE) has become the industry standard for measuring the energy efficiency of data center facilities. PUE was established in 2007 by a consortium called the Green Grid, and has been an ISO standard since 2016.[1]

Before the widespread use of PUE, there was no comparable metric for analyzing the performance of data centers over time, or judging the relative power consumption and distribution of a given facility as it relates to its efficiency.

PUE is a very helpful tool. But, like almost any metric used in isolation, it has its shortcomings. There are two primary issues with the PUE metric: first, the misleading nature of factors used in its calculation, and second, the factors missing from that calculation.

In this paper, we'll provide an overview of PUE, including what impacts it and how to optimize it, and then dive into what is missing from the PUE calculation. We will conclude with recommendations for other industry-wide standards that should become common practice.

What is Power Usage Effectiveness?

PUE measures the relationship between total energy consumed and information technology (IT) equipment energy consumed in a given facility.

$$PUE = \frac{(\text{Total Facility Energy Consumed})}{(\text{IT Equipment Energy Consumed})}$$

PUE provides a means to determine:[2]

- Opportunities to improve a data center's operational efficiency.
- How a data center compares with similar data centers.
- If data center operators are improving designs and processes over time.
- Opportunities to repurpose energy for additional IT equipment.
- A design target or goal for a new data center.

IT equipment energy is "the energy consumed by equipment that is used to manage, process, store or route data within the compute space."[3] IT equipment includes computing hardware and storage and network equipment, as well as monitors, switches and workstations used to monitor or control the facility. These are all components included in the denominator of the equation.

The numerator of the PUE calculation takes into account all IT equipment as well as all non-IT equipment that uses energy inside the facility. This includes power delivery components, including UPS systems, switchgear, generators, power distribution units (PDUs) and batteries. Energy is also consumed by distribution losses external to the IT equipment, cooling system components (chillers, cooling towers, pumps, computer room air handling units, computer room air conditioning units and direct expansion air handler units) and other miscellaneous component loads, such as data center lighting.

A ‘Good’ PUE

PUE is expressed as a ratio. Overall efficiency improves as the quotient decreases toward 1.

A PUE of 1.0 means for every 100 megawatts of energy going to a site, 100 megawatts go to the equipment that is used to manage, process, store or route data within the compute space. A PUE of 1.0 is perfect (and impossible to achieve).

There is no broadly accepted gold standard for PUE. NREL cites the average PUE as 1.8, and notes that “data centers focusing on efficiency typically achieve PUE values of 1.2 or less.”[4] Another way to think of a PUE of 1.8 is that for every 100 MW that go to the facility, 55.55 MW go toward the computing needs in the facility.

$$1.8 = \frac{100}{55.55}$$

Google’s best data center in 2021 achieved a PUE of 1.06.[5] Soluna’s first data center has achieved a consistent PUE of 1.02.

Google / PUE of 1.06

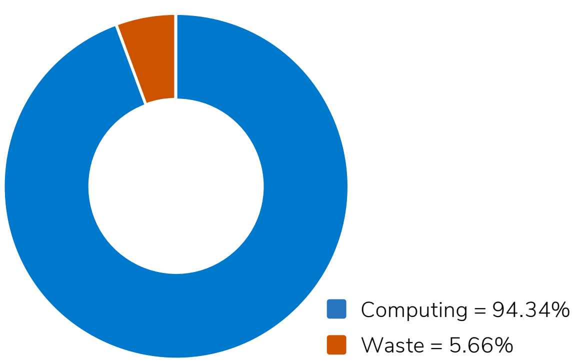

Soluna / PUE of 1.02

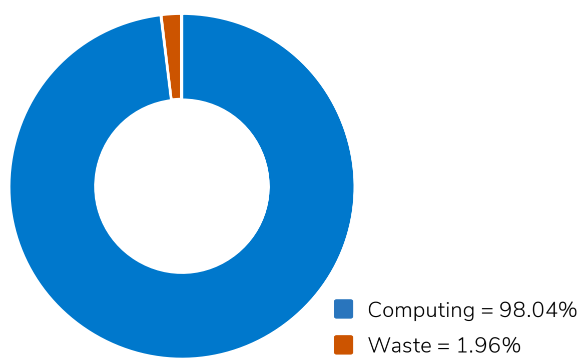

What Impacts PUE?

Several factors impact PUE, including the type of work done in the facility, the physical structure, the design and operation of the facility, and where the facility is located. There are several parameters to each of these factors, as follows:

A	**Type of work done in the data center** *(e.g., testing, production, internal processes, networking)*	• Primary business supported by the data center. • Level of resiliency required to support business operations. • Bottom line: Different types of work use energy in different ways.
B	**Physical Structure**	• **Temperature and humidity levels.** Natural temperature and humidity impact the state of equipment and ancillary cooling system needed. • **Cooling system used.** The mere existence of a cooling system uses significant energy. • **Age of data center and subcomponents.** Older equipment might not use energy as efficiently. • **Building's intended use.** Some buildings aren't intended to be retrofitted as data centers, and all subcomponents might not be relevant to the end-use of computing. • **Energy-reduction features.** Site design and ancillary equipment can significantly impact a site's energy usage.
C	**Design & Operation**	• **Power distribution architecture.** The way that power sources are designed and configured throughout the facility. • **Redundancy levels.** A system design where a component is duplicated so that in the event of power outage, equipment failure or unexpected breakage, IT equipment will not be affected in an alarming way. Redundancy levels are measured via tiers.
D	**Climate & Location**	• **Local climate.** Might affect overall efficiency and hours of free cooling available; a colder climate means less need for cooling systems. • **Location.** Might determine access to certain types of equipment, supply chain availability, etc.

The cost and climate impact of IT is unsustainable on its current trajectory.

Achieving Near-Perfect PUE

Soluna utilizes a unique proprietary site design. Our sites are constructed modularly, like LEGO® blocks. Our buildings are arranged in a diamond configuration, with four buildings making up each individual diamond.

fig 4.1: Diamond configuration at Soluna's Project Sophie facility

Several factors contribute to our low PUE and high operational efficiency:

A **Cooling**

Cooling technology uses a lot of energy, and allowing server rooms to operate at a higher temperature or taking advantage of lower outside temperatures can dramatically reduce energy usage.

Efficient thermodynamics is key here. Soluna modular data centers (MDCs) are cooled by nothing more than very efficient airflow. Currently, our buildings all house cryptocurrency miners. Miners are way more energy dense than typical servers, so the conditions we have that keep our miners cool will keep any other hardware cool as well.

B **Airflow**

Airflow is primarily determined by fan speed, which is influenced by fan size. Bigger fans can turn slower to get the same amount of airflow. We utilize giant fans on the sides of our buildings that operate at an RPM of 300 (a typical desk fan has an RPM of 1,200). The relationship between fan speed and power draw isn't linear. A fan that runs at 1/4 speed draws only 1/8 as much power. There are also efficiencies in having a larger diameter. Overall, the big fans on the sides of our buildings are around 20 times as efficient at moving a given volume of air as the little fans that cool our hardware. They also produce significantly lower sound levels.

Our sites are designed so that the wind will help us regardless of the direction it's blowing. If prevailing winds are blowing into the exhaust, the entire thermal system has to work harder to move the same amount of air. Our building exhausts face inward, toward one another, so all of the air being produced by the fans blows toward the center of the site. This means that regardless of the prevailing wind speed and headings, our buildings are protected from these headwinds.

Our airflow is designed to minimize bends and blockages. Air goes through the filters, into the miners, and out the back of the facility. We move the air slowly, never move it through a narrow duct, and keep our facilities from recycling air (which reduces efficiency). Both tilted roofs and the diamond shape of our facilities serve to reduce recirculation.

We deploy filters in extended louvers instead of installing air filters directly onto the sides of buildings. The surface area of the filters is twice that of the intake wall of the buildings. More filter area means less flow restriction as the air moves through the filters and less work for the fans. This configuration also protects the filters from rain and other debris.

C

Other aspects of site design

We don't have any ancillary equipment beyond network equipment and storage. Minimal ancillary equipment means more of the energy being produced goes directly to computing.

Additionally, we take advantage of expansive sites that allow us to spread things apart. We're able to do this because our sites are location-agnostic, and can be located where land is plentiful. Unlike traditional data centers, we don't require a convergence of strong bandwidth and broadband, which allows us to set up our sites in sparsely populated areas where there is an abundance of wind or solar power — and an abundance of space.

Renewable Computing:

a new class of 011010110 infrastructure

Issues with the Way PUE is Measured

The official PUE metric includes all IT equipment in the denominator. We consider this problematic, as the purpose of calculating PUE is ascertaining what percent of the energy going into your facility has a productive use (i.e., is going directly to computing).

IT equipment is still supporting infrastructure. (Much like storage or lighting, IT equipment supports the end-use of computing, but is not itself the end-use of energy entering the facility.) Including IT equipment in the denominator of the calculation obfuscates the true cost of inefficient networks and sloppy engineering.

This also means that while PUE provides a measure of energy distribution in a facility, it doesn't actually provide insight into the usefulness of that facility.

A low PUE calculation that includes IT equipment in the denominator might still have inefficient IT equipment and thus low output.

At Soluna, we address this ambiguity by including only the power that goes directly to computing (only the energy that powers computing hardware, e.g., miners) in the denominator. We recommend that the PUE metric be revised accordingly, thus providing a stricter metric that demands more from computing enterprises.

<u>*fig 5.1:*</u> PUE Calculation Comparison

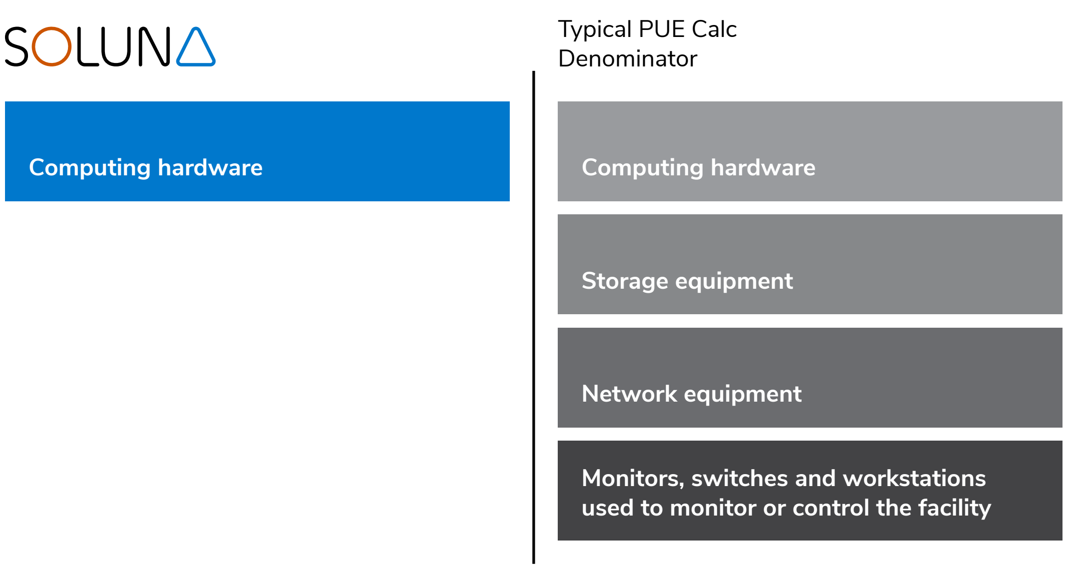

What is Missing in the PUE Calculation

PUE is a great starting point, but it was never intended as a stand-alone metric and shouldn't be utilized as one. In order to capture a more complete picture of data center performance, we recommend including the following factors in assessments:

A **Efficiency of chips**

If you're using old or inefficient hardware, you're wasting energy because of low relative output. PUE measures only the input side of the equation — how much energy is going to IT equipment versus other purposes. Without a clear understanding of the relationship between input and output, you cannot effectively measure efficiency of a data center.

B **Maintenance of equipment**

During operations, ancillary pieces of hardware (power supplies, onboard cooling, and ancillary control units) are the leading cause of failure. Soluna places an emphasis on better data center design, which not only reduces energy consumption but also reduces equipment failures.

C **Volume of water being used**

A typical data center uses 3 to 5 million gallons of water per day.[6] Pure, potable water that could serve a city of 30,000-50,000 people. Soluna data centers use no water for cooling; they are air-cooled. Our data centers operate with world-leading efficiency, and they do so without consuming any water.

D **Location**

Many data centers are located in or close to urban hubs that provide both high bandwidth and a strong internet backbone. Centrally located data centers drive up rent, contribute to traffic congestion, and produce noise pollution for the surrounding area.

E **Waste Materials**

PUE doesn't take into account electronic waste ("e-waste"), which presents a huge challenge to environmental sustainability. E-waste is a concern for all businesses operating in industries that rely on information technology equipment. Other solutions to data center efficiency like immersion cooling waste other resources (e.g., oil).

F **Source of power**

PUE doesn't take into account what kind of energy you're using and its carbon footprint. It's quite possible to have a low PUE in a facility that is powered by coal, which undermines any claims regarding energy efficiency.

Conclusion

Since its inception in 2007, PUE has dramatically changed the landscape for assessing data center performance. The issues with PUE as laid out in this paper are two-fold.

If the primary goal is optimizing facility efficiency, PUE should calculate computing processes against all other supporting infrastructure (including IT equipment). This revision in the formulation of PUE would provide a more accurate picture of the real usefulness of a facility.

But the larger point is that as data centers proliferate, we should be demanding more from them. The global cloud computing industry grew by 13.7% in 2020 alone,[7] and is projected to hit $791.5 billion by 2028.[8] Energy demands have not increased commensurately, thanks to a growing focus on sustainability and climate-first efforts. From 2010 to 2018, data center computing workloads increased almost 550%, but data center electricity consumption rose only 6%.[9]

But the pace of cloud computing is only accelerating. According to a letter published in the open-access scientific journal Environmental Research Letters in May 2021, "The amount of data created and stored globally is expected to reach 175 Zettabytes by 2025, representing nearly a six-fold increase from 2018."[10]

PUE was not intended to capture the full range of externalities associated with data centers. These include carbon footprint, as well as myriad other factors not taken into account when assessing data center performance.

There are clear benefits to distilling a number of complicated factors into a single metric to communicate performance. If PUE is to be that metric, it should be calculated more stringently *and* factor in a more complete picture of what efficiency actually means. This single metric could be further substantiated by other factors, including the maintenance of equipment and e-waste practices. But at minimum, a truly useful metric for assessing data center performance would calculate the real usefulness of a facility, as well as its associated carbon footprint.

The Soluna Team

We're Solving Renewable Energy's Biggest Problem.

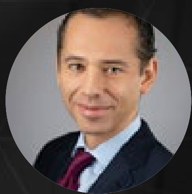

Michael Toporek
Chief Executive Officer
Soluna Holdings, Inc.

John Belizaire
Chief Executive Officer
Soluna Computing, Inc.

Philip Patman
Chief Financial Officer
Soluna Holdings, Inc.

Dipul Patel
Chief Technology Officer
Soluna Computing, Inc.

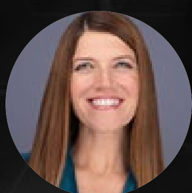

Mary O'Reilly
Chief People Officer
Soluna Computing, Inc.

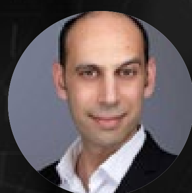

Larbi Loudiyi
VP, Power
Soluna Computing, Inc.

Phillip Ng
VP, Corporate Development
Soluna Computing, Inc.

Nick Lancaster
VP, Engineering
Soluna Computing, Inc.

Chase Stowell
Controller
Soluna Holdings, Inc.

Anna-Lynn Mancini
VP, Operations
Soluna Computing, Inc.

Sam Sova
VP, Marketing
Soluna Computing, Inc.

To learn more about how Soluna is bringing renewable computing to the grid, join our community here.

Reference Notes

1 ISO
https://www.iso.org/standard/63451.html
accessed September 20, 2022

2 The Green Grid
https://datacenters.lbl.gov/sites/default/files/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric_v6.pdf
accessed September 15, 2022

3 The Green Grid
https://datacenters.lbl.gov/sites/default/files/WP49-PUE%20A%20Comprehensive%20Examination%20of%20the%20Metric_v6.pdf
accessed September 15, 2022

4 NREL
https://www.nrel.gov/computational-science/measuring-efficiency-pue.html
accessed September 22, 2022

5 Google
https://www.google.com/about/datacenters/efficiency/
accessed September 12, 2022

6 NBC News
https://www.nbcnews.com/tech/internet/drought-stricken-communities-push-back-against-data-centers-n1271344
accessed September 27, 2022

7 Nasdaq
https://www.nasdaq.com/articles/cloud-computing-energys-efficiency-problems
accessed September 27, 2022

8 IBID

9 IOP Science
https://iopscience.iop.org/article/10.1088/1748-9326/abfba1
accessed September 27, 2022

10 IBID