

Energy Efficiency in Data Centers: Moving Beyond PUE

Nathaniel Horner & Inês Azevedo

Department of Engineering and Public Policy, Carnegie Mellon University

PUE is the standard industry metric for measuring data center efficiency.

It is a measure of the energy overhead needed to operate a data center:

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

Energy consumed by the memory, chips, and hard drive goes in the denominator, while the numerator includes this energy plus energy used for cooling, lighting, and other support functions.

The energy usage breakdown for a typical data center shows that only about half of the load is used directly for computing.

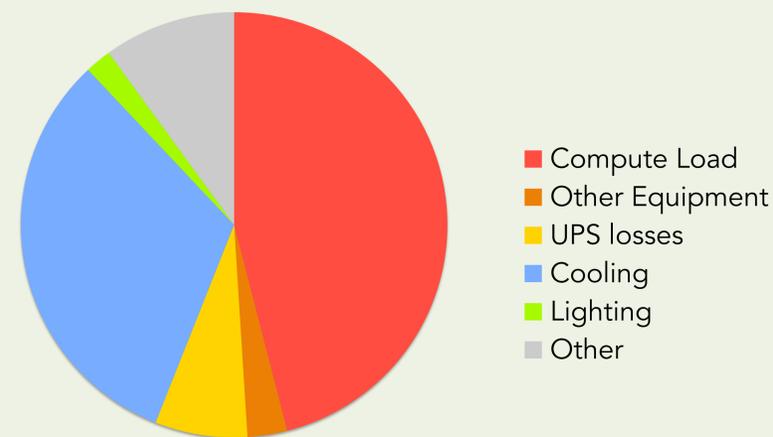


Fig. 1: Energy usage in a typical data center.

The data center in Fig. 1 has a PUE of 2.17. Google and Facebook report PUEs as low as 1.08. The estimated national fleet average is 1.8 – 2.9. [1,2]

[1] Stansberry M and Kudritzki J 2012 2012 Data Center Industry Survey (Santa Fe, NM: Uptime Institute)
 [2] Digital Realty Trust 2013 North America Campos Survey Results (San Francisco, CA: Digital Realty Trust)
 [3] Most data from LBNL 2003 Benchmarking: Data Centers – Case Study Reports High-Performance Buildings for High-Tech Industries and from Rod Mahdavi at LBNL, supplemented with publicly-reported information on two federal and four private data centers.

But PUE is not necessarily correlated with energy or emissions performance.

Data from a set of 40 facilities (Fig. 2) shows that low PUE values do not imply low energy or carbon intensity.

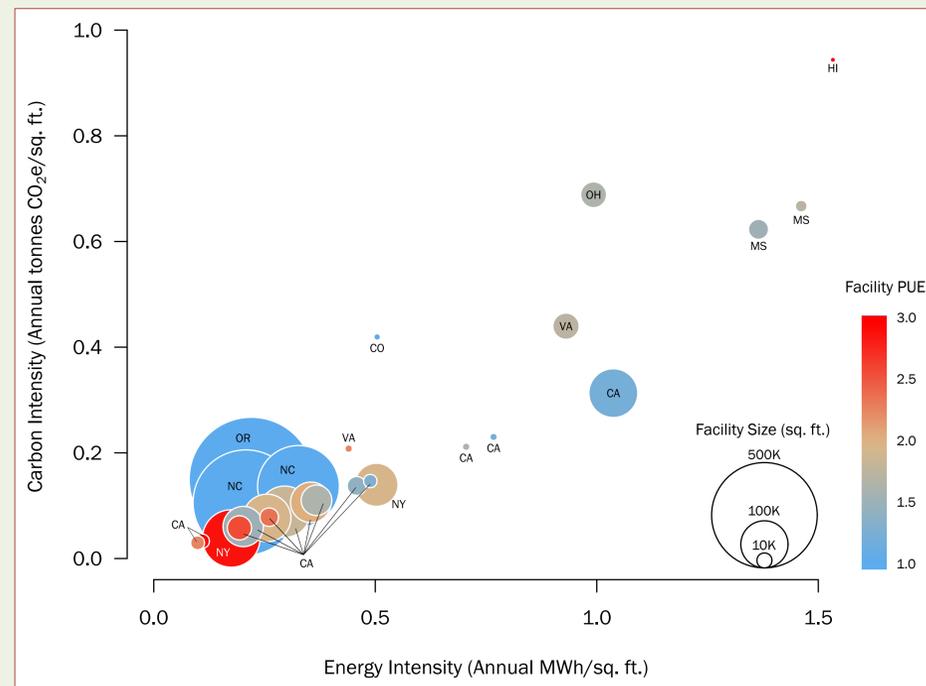


Fig 2: Annual energy intensity (MWh/sq. ft.) vs. GHG emissions intensity (t-CO2e/sq. ft.) for 26 U.S. data centers [3] represented by facility size and PUE rating. (Additional data centers omitted for readability.)

See, for example, the CO and NY facilities with energy intensity of 0.5 MWh/sq. ft. Despite having a better PUE, the CO data center has much worse carbon intensity.

This happens because PUE does not account for electricity source—and the metric misses many other data center characteristics, too, such as the IT efficiency. Perversely, a data center that upgraded to more efficient hardware would likely see its PUE increase.

[4] Natural Resources Defense Council 2012 The Carbon Emissions of Server Computing for Small- to Medium-Sized Organizations: A Performance Study of On-Premise vs. The Cloud
 [5] Barroso L A and Hölzle U 2009 The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines Synthesis Lectures on Computer Architecture 4 1 pp 1–108
 [6] Patterson M K, Poole S W, Chung-Hsing H, Maxwell D, Tschudi W, Coles H, Martinez D J and Bates N 2013. TUE, a new energy-efficiency metric applied at ORNL's Jaguar Supercomputing 7905 pp 372–382

Further, PUE says nothing about the productivity of energy in a data center.

Efficient cloud data centers have 40% utilization rates. [4] The worst data centers are in the single-digits. Because servers are not energy-proportional (they use a significant amount of power when sitting idle), low utilization means that energy consumed by computing hardware is not being used for IT service delivery.

We need other metrics.

The problem is not with PUE itself, but rather that the industry has gotten caught up in a “PUE race.” It’s time bring other metrics into the picture.

In addition to PUE at the meter, we could also measure IT overhead at the server box: SPUE [5] or ITUE [6] is the ratio of energy into the server box to energy translated into computing work by the individual components.

Carbon Usage Effectiveness (CUE) [7] and Energy Reuse Effectiveness (ERE) [8] have been proposed by The Green Grid as green-focused metrics.

Utilization is perhaps the most promising area of focus, as increasing utilization rates would seem to be an obvious, large improvement. However, there is no standard for how to measure utilization. We are in the initial stages of a project with Lawrence Berkeley National Laboratory to define such a standard.

[7] Belady C 2010 Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric (Beaverton, OR: The Green Grid)
 [8] Tschudi, B. et al. 2010 ERE: A metric for measuring the benefit of reuse energy from a data center (Beaverton, OR: The Green Grid)

Images of Google Data Centers
<http://www.google.com/about/datacenters/>

