Prism Power Group

Switchgear • Critical Systems • Intelligent Solutions • Maintenance Services

Data Centre Efficiency and the Limitations of PuE

Ian F Bitterlin
PhD BSc(Hons) DipDesInn MCIBSE MIET MBCS MBIFM
Chief Technical Officer
Prism Power Ltd, Watford, UK

Prism Power Limited delivers high integrity Critical Power Solutions.

This document is one in a series of white papers prepared by Prism Power. If you would like further information, please contact Prism Power on +44 (0)1923 296 700 or visit www.prismpower.co.uk

Introduction

Even before electrical energy costs started to climb in 2001 data-centre designers were thinking about the relationship between the power consumed by the IT hardware itself and the total power absorbed by the containment facility. The driver for that process in the dot.com boom was the apparent shortage of power capacity in various internet-hub locations such as London, Dublin, Amsterdam, Paris and Frankfurt and the desire to maximise the power available for the IT hardware. The thrust of the thinking then, as now, was focussed on the Coefficient of Performance[1] of the cooling system and the power required to recharge the UPS after an outage.

At that time the IT hardware costs and the facility construction costs were dominant and the concentration was on the space efficiency; the ratio of raised floor area to total built area.

However, since the dot.com crash in 2001 key drivers have influenced the discussions around 'efficiency':

- Construction costs have only risen with inflation
- IT hardware costs have fallen dramatically per unit of performance
- Energy costs have generally doubled and are continuing to rise far ahead of inflation
- Grid capacity has become more difficult to source
- Corporate entities have adopted 'green' as a mantra
- Data-centre vendors all want to parade their green-credentials

Power costs have become the largest item in the OpEx budgets of data-centre owners and, for the first time, now exceed the IT hardware cost over its average service-life of 3-4 years. This has resulted in the pursuit (both real and spun) for higher efficiency, sustainable, 'green' designs.

PuE, DCiE et al

The basis of most improvement programmes involves the old adage *'you can't control what you don't measure'* so the development of a common measurement tool for data-centre energy efficiency has been high on many agendas. Despite their various heritages[2] all of the definitions come down to some form of the relationship between IT power and total site power taken from the electrical grid.
The most commonly accepted definitions are the PuE (Per Unit Efficiency) and the DCiE (Data Centre Index of Efficiency):

• PuE is the Site power divided by the IT power. E.g. Typically ranges from a 'good' 1.6 to a 'poor' 3 – so that 1 is (unobtainable) 'perfection' and the number gets larger as the efficiency gets worse
• DCiE is the inverse of PuE, therefore the output ranges from 0.6 to 0.3 – with the numerical result getting smaller as the efficiency declines

1 The mechanical cooling system does the job of capturing the heat generated by the IT equipment and transferring it to the outside for rejection.
2 The Coefficient of Performance (CoP) expresses the efficiency of that process. E.g. If it takes 400W of power to remove 1000W of heat the CoP is 0.4 or 40%

There is still debate about how the two input power are measured (and where) but there is general agreement that the method is valid and valuable if carried out on a single facility to monitor progress between phases of improvement.

For example a particular facility might be measured for IT power (the sum of all the PDU[3] inputs) and total mains power just before a phased programme of works such as:

- Blanking plates, hole-stopping & remove redundant under-floor cables to manage air-flow
- Hot-aisle/cold-aisle layout implementation with cold-aisle (or hot-aisle) containment
- Raised room set-point temperature with raised chilled-water temperatures
- Implementation of free-cooling equipment

At the end of each work-phase a re-measurement and resultant new PuE calculation will be a very useful exercise – indeed it will also provide data for the facilities EU Energy Efficiency rating. In this respect (using PuE or DCiE) as a measurement of relative improvement the process is a relevant and valuable tool.

There are however some basic flaws even at the lowest level:

- A data-centre with a large office (for Disaster Recovery for example) as part of the facility will have any efficiency gains within the data-centre masked in the roll-up of the site numbers and appear to have a high PuE (poor) unless the data-centre load can be measured in isolation
- A data-centre which uses other forms of energy (other than mains electricity) will appear to have a very good PuE. For example:
    - Cooling systems powered by natural gas (absorption chillers) can have a PuE below 1.2 because nowhere in the PuE calculation is the energy content of the natural-gas
    - On-site photovoltaic solar panels (for small facilities)
    - On-site production of electricity & heat via co-generation, particularly using sustainable fuels such as bio-diesel, wood-chip, straw or municipal waste
- A data-centre with free-cooling installed can be measured in winter and then compared with summer a measurement – probably showing a marked improvement
- A lightly loaded facility (perhaps new) will appear to have a very poor PuE which will increase as the load grows – with no improvement action on the part of the operator

The measurement and calculation of PuE is in the hands of the facility owner/operator and, so far, there is no regulatory certification body, and so, just like Income Tax, self-certification is open to individual interpretation.

Black-box thinking?

Having established that the use of PuE is a valid (and valuable) tool for measuring progress in a particular facility we have now to turn to its principle 'flaw'.

3 Power Distribution Units in the data-room

For a corporate entity the whole issue of energy efficiency can be viewed two ways:

• Improving energy efficiency saves money and happens to be 'green' – and, as long as it does save money, we want to be 'green'
• We want to be 'green', even at a premium capital cost

So far, in the UK and USA at least, most businesses fall into the first category – even to the extent that a Return-on-Investment of longer than 2 years for higher cost capital plant will often be rejected. However, there is a growing desire to appear green, driven by cost saving and, as the cost of energy spirals upwards, this will increase.

But what does 'green' really mean? It means sustainability – in particular, for power, it means no (or lower) carbon footprint. However there are three steps to sustainability that have to be taken in precise order and efficiency is the second step, not the first:

1. Reduce the consumption of the process
2. Improve the efficiency of the process
3. Source the energy from sustainable (renewable) sources

In this way it can be seen that building an inefficient data-processing facility and powering it from local hydro-power is not truly sustainable – it is better considered to be a waste of a valuable carbon-free resource.

So, to be truly 'green' and sustainable we need to address the matter of the load consumption before we worry about the efficiency. It is not within the scope of this paper or the electrical system designer to be concerned about the value and content of the data being processed – that is far more philosophical discussion about modern society – so we have to accept that the electrical input to the load is our starting point.

The power consumed by a micro-processor based load falls into three segments

• The power consumed by the core of the micro-processor, typically 70-130W each
• The power dissipated by the on-board power supply, typically 30-50W per processor
• The power required to drive the cooling fans to enable the heat to be rejected, typically 20-30W per processor

In this way we can see that each core consumes between 120-210W[4]. Add all the processors up in a fully populated blade server chassis (often >30 cores), along with the usual hard-drive motors and controllers, and it is easy to reach 5kW for a server in a 7U space.

Of course, every Watt that enters the load terminals is rejected as heat and so it is possible to criticise the concept of PuE because it ignores the fundamental 'efficiency' of the load itself. It would be more logical (but extremely difficult) when calculating the PuE to consider the power consumed by the cores and take the power-supply losses and fan power as part of the site-power – in that way the pressure for high efficiency would be extended to the switched-mode power-supply and the internal cooling. That said,

4 Also has the undesirable characteristic of consuming >70% of the maximum power when inactive

it could also be considered that the performance of the micro-processor itself (in terms of instructions/second/Watt etc) is a fundamental part of the evaluation process.

The immediate conclusion is that PuE ignores the efficiency of the load, but shouldn't. Since low-power cores are available, along with high-efficiency power-supplies and variable speed low-power fans, this omission is fundamental. The client should start with low-power high-efficiency servers and then proceed to considering the facility's PuE – but he/she will never be praised for the substantially lower power consumption of their data-centre because the PuE will ignore that extra effort and, usually, investment.

An holistic view of energy consumption

As we have seen in the previous section it is not just the part of the electrical energy stack that the PuE methodology covers that matters – not forgetting the input of other fuels like natural-gas for chillers etc. The diagram, below, illustrates where the majority of the losses occur in a data centre.

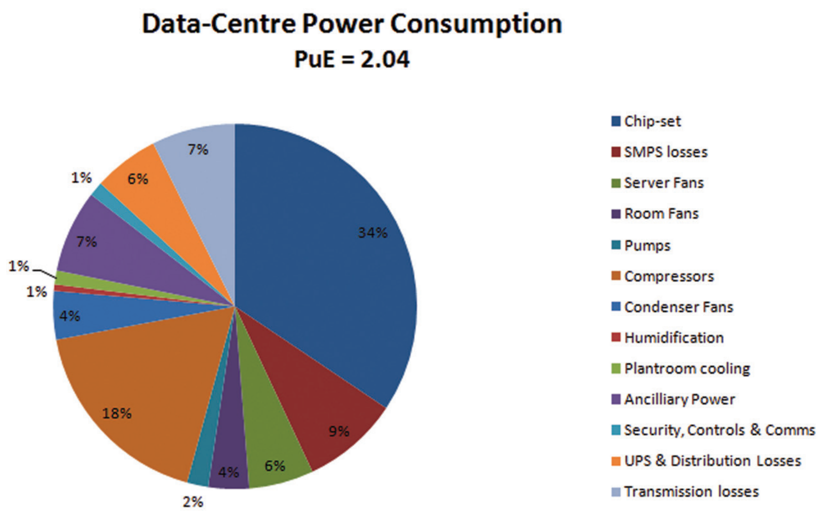

Whilst the PuE covers only 54% of the whole energy stack between power-station and IT load terminals we should be aware of the elements that can vary widely and dramatically affect the overall picture:

• The chip-set itself; with low-power chips now available and more forecast down to 40% of current power-draw
• The on-board switched-mode power-supply; improving from 60% to 85% efficient
• The mechanical cooling system; here shown totalling 30% (CoP = 0.6) but often as high as 50% for poorly designed air-flow and can be as lower than 20% (CoP = 0.4)
Whilst it is good that the PuE captures the mechanical cooling system the load itself can be halved by careful selection and specification of hardware.

Redundancy and the problems of partial load

It will be of no surprise to the engineering designer that it is not possible to construct commercially attractive plant that does not have peak efficiency at a particular load and this certainly applies to data-centre plant such as UPS, chillers, room air-conditioning units, IT hardware power-supplies etc. Most plant is designed to run near peak efficiency usually in the upper quartile of load.

However data-centres have two features that result in low loads on the critical equipment:

- They are rarely fully loaded for the majority of their service lives
- They are designed for reliability with high levels of redundancy built-in to avoid down-time

The combination of both factors often leads to loads on plant of less than 15% - where no plant is energy efficient. For example take a Tier IV UPS system with A+B paths and 2+1 redundancy in each path. At a room load of 40% each UPS module runs at less than 14%.

Compare this to a Tier I system (single path with no redundancy) of the same capacity – the UPS runs at 40% load, not an ideal set-point but at least 20%age points better in operating efficiency.

Maybe unsurprisingly are we to infer that high-reliability (via redundancy) is the enemy of high-efficiency and , hence, of PuE?

Conclusion: There is danger in PuE becoming a marketing tool for facilities

Like the exaggerated claims for Tier Classification (all self-certified) the claims for PuE (or DCiE) have started to appear in sales-pitches for data-centre space. So the following caveats should be considered:

- These claims are made for centres not yet built and must assume that the centres will be >75% occupied from the day the doors open. The PuE of a new centre will be poor until the systems are loaded up
- Using PuE as a yard-stick to compare different locations is highly dubious
- You cannot compare the PuE of a Tier I facility with a Tier IV (or Tier II or III). If you do then you will accept that energy efficiency overrides reliability and availability in your data-centre
- It is a valid tool for recording improvements in one centre – but even then is open to abuse and misinterpretation
- Always start with selecting energy-efficient IT hardware with low-power micro-processors

Prism Power Group

Switchgear • Critical Systems • Intelligent Solutions • Maintenance Services

Ian F Bitterlin
PhD BSc(Hons) DipDesInn MCIBSE MIET MBCS MBIFM
Chief Technical Officer
Prism Power Ltd, Watford, UK
e. ian.bitterlin@prismpower.co.uk