Ark Continuity - Continuously Delivering high integrity data centres

The Holy Grail of “Green” Data Centres... is not PUE

Dr Ian F Bitterlin

PhD BSc(Hons) BA DipDesInn MCIBSE MIET MBCS MBIFM MIEEE

Publication: September 2010 | **Reference** 2010/09/006

WHITE PAPER 2010

Ark Continuity delivers high integrity data centres for Government and Corporate occupiers. Ark is now building data centre campus locations at Spring Park, Wiltshire and Cody Park, Hampshire. Data Centre SQ17, module one, is now operational. All facilities are designed to be the most secure, available and sustainable in Europe.

This document is one in a series of white papers prepared by Ark Continuity. If you would like further information, please contact Ark Continuity on 0845 389 3355 or visit www.arkcontinuity.co.uk

Head office: Hartham Park, Corsham, Wiltshire, SN13 0RP

The Holy Grail of "Green" Data Centres

...is not PUE

Introduction

Even before electrical energy costs started to accelerate in 2001 data-centre designers were thinking about the relationship between the power consumed by the IT hardware and the total power absorbed by the facility. The driver for that process in the dot.com boom was the shortage of grid 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 was focussed on the Coefficient of Performance[1] of the cooling system and the power required to re-charge the UPS after a grid power outage.

Before 2001 the IT hardware and facility construction costs were dominant and the design 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 risen in line with inflation
- IT hardware costs have fallen dramatically per unit of performance
- Energy costs have generally doubled and are forecast to rise ahead of general inflation
- Grid power capacity has become even more difficult to source
- Corporate entities have adopted 'green' as a mantra and data-centre vendors all want to parade their green-credentials

Power costs have become the largest item in the OpEx budgets of data-centre operators and, since 2008, 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 (Power Usage Efficiency) and the DCiE (Data Centre Index of Efficiency):

---

[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. 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%

[2] Including the work of both independent and industry led organisations; such as The Green Grid, British Computer Society, The Carbon Trust, The Uptime Institute, Energy Star Programme etc - even the 2009 EU Code of Conduct on Data-Centres

- 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

There is still debate about how the two points are measured (and where) but there is general agreement that the method is a valid and valuable tool 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 improve 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 in the concept of PUE 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 IT load and facility overhead 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 very 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 – always showing a marked improvement
- A lightly loaded facility (e.g. new) will have a very poor PUE which will improve dramatically as the IT load grows – despite 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.

[3] Power Distribution Units in the data-room

Black-box thinking or sweeping the problem under the carpet?

Having established that the use of PUE is a valid and valuable tool for measuring progress in an individual facility we have now to turn to its principle 'flaw'.

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 an ROI (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 a low(er) 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 the M&E engineer has 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. Add all the processors up in a fully populated blade server chassis (often >30 cores), along with the usual hard-drives, controllers and I/O, and it is easy to reach 5-6kW for a blade chassis 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 infrastructure-power – in that way the pressure for high efficiency would be extended to the switched-mode power-supply and the internal cooling. That said, 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. In fact, reducing the actual load will usually increase the PUE (i.e. negatively).

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, right, illustrates where the majority of the losses occur in a data centre.

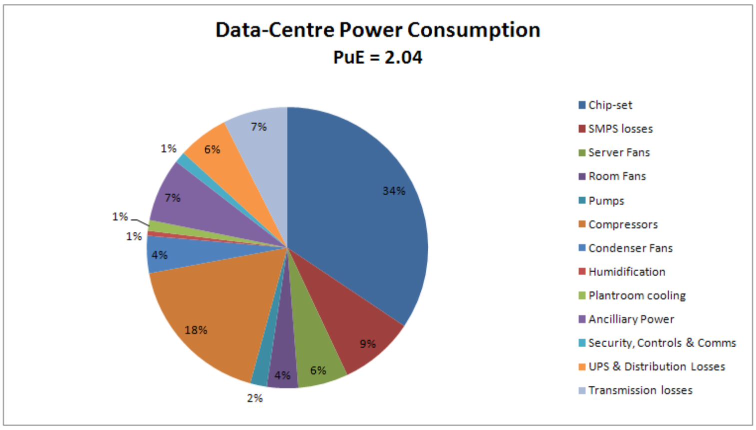

Whilst the PUE covers only 54% of the whole energy stack between power-station and IT load terminals we should be aware of the elements that can vary widely and dramatically affect the overall picture:

- 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 90% 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 IT load itself can be drastically reduced by additional expenditure and careful specification of IT hardware.

So is PUE the Holy Grail?

Despite all of the pressure brought to bear on PUE the answer to this question is a definite 'NO'.

Data-centre systems have improved dramatically, from a PUE of well over 2.5 down to 1.5, especially in Northern climes where free-cooling is abundant. When all of the recommended best practice air-management measures are taken[4] and free-cooling is available, a 'coldest day' PUE of 1.3 is relatively easy to achieve at minimal additional investment. However, if fresh-air is introduced into the critical space then PUEs in the region of 1.15 are more than possible. The energy cost saving makes this option increasingly attractive. When you consider the two extremes, 2.5 to 1.15; an annual saving[5] of GB£1.18M per MW of IT load is a very compelling argument in favour of fresh-air solutions.

So where is the Holy Grail if it is not PUE? It is inside the IT hardware 'box'. Although we have already considered the losses in power conversion, cooling fan, hard-drive and I/O power we should look at the relationship between processor 'activity' (the nearest thing we can get to a definition of 'useful work'?) and the power consumed by the server.

The problem is that the average server consumes >60% of its peak power when doing no processing work whatsoever. The input power load-line is approximately linear between zero compute load and 'flat-out'. The graph, right, shows this simple relationship in the average pizza-box server and the red-dotted line illustrates where the 'curve' should lie.

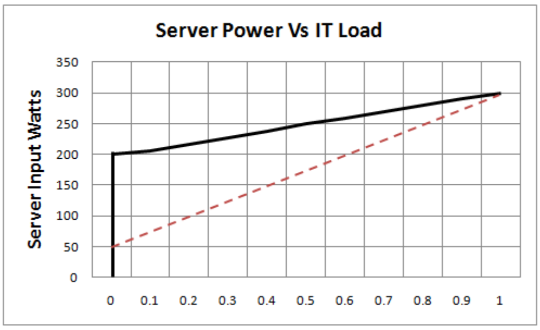

The problem of base-load power draw is compounded by the fact (it is universally accepted) that the average server microprocessor chip-set is only 5% utilised – and certainly rarely higher than 10% utilisation. That goes a long way to explain why server power is never measured to be anywhere near its name-plate data. A fully configured blade chassis will usually draw 60-65% of what is expected. This '60%' base/peak performance characteristic is not affected by high-efficiency power supplies, solid-state drives or low power chips or I/O – only the overall power reduces – and, due to the incredibly high compute capacity and speed of the modern chips, virtualisation hardly improves the chip utilisation situation, so the fact remains that the IT 'zero work' base load is 60% of the whole facilities power demand regardless of the PUE.

Conclusion

To be truly sustainable the consumption of any process must be minimised and '60%-input for 0%-output' (less than 0% efficiency?) must be a good place to start that process.

Whilst M&E engineers beat themselves up to gain improvements of 0.1 (e.g. from 1.6 to 1.5) and the data-centre industry faces penalisation under the CRC for lack of year-on-year energy reduction the real area for investigation and improvement is the zero-compute power demand of the IT load itself.

End

To read more white papers from Dr Ian Bitterlin, please visit www.arkcontinuity.co.uk

[4] Physical aisle-containment, blanking plates, hole stopping, higher server inlet temperatures etc
[5] Calculated at GB£0.10/kWh