

Sustainable Data Centre

Tackling the Ever-rising Data Centre Energy Costs – Why Only Do Half the Job?

By Kari Baden – Managing Director, Dimension Data Advanced Infrastructure



A whitepaper on tackling the ever-rising data centre energy costs.

Abstract

Businesses today couldn't operate without IT and to many organisations this means building and running a data centre. This brings with it a whole new set of challenges which have evolved over time and are now focused on raising energy supply costs and energy security. This has come about because of the rapid growth of IT infrastructure and the rapid increases in energy provision. Energy simply didn't register on the data centre radar, but now is a cost line that must be addressed. There are a range of approaches that can be used to tackle the challenges specific to each client. It's important to consider the whole range of options to build your plans to tackle the situation and to get started.

Introduction

For those with a responsibility for running a data centre, there's most likely already pressure to deliver the availability of applications to internal and external users and to prevent security attacks. On top of these traditional pressures, there's now also the task of managing escalating energy costs. The good news is that there are a range of ways to address the energy question; the challenge is that many organisations are only presented with half the answers by suppliers in the market. The purpose of this whitepaper is to take a holistic view of the solutions available, in order to help those responsible for building, operating and managing data centres determine the options that are best for their organisation.

Energy simply didn't register on the data centre radar, but now is a **cost line** that **must be addressed**.

How did we get here?

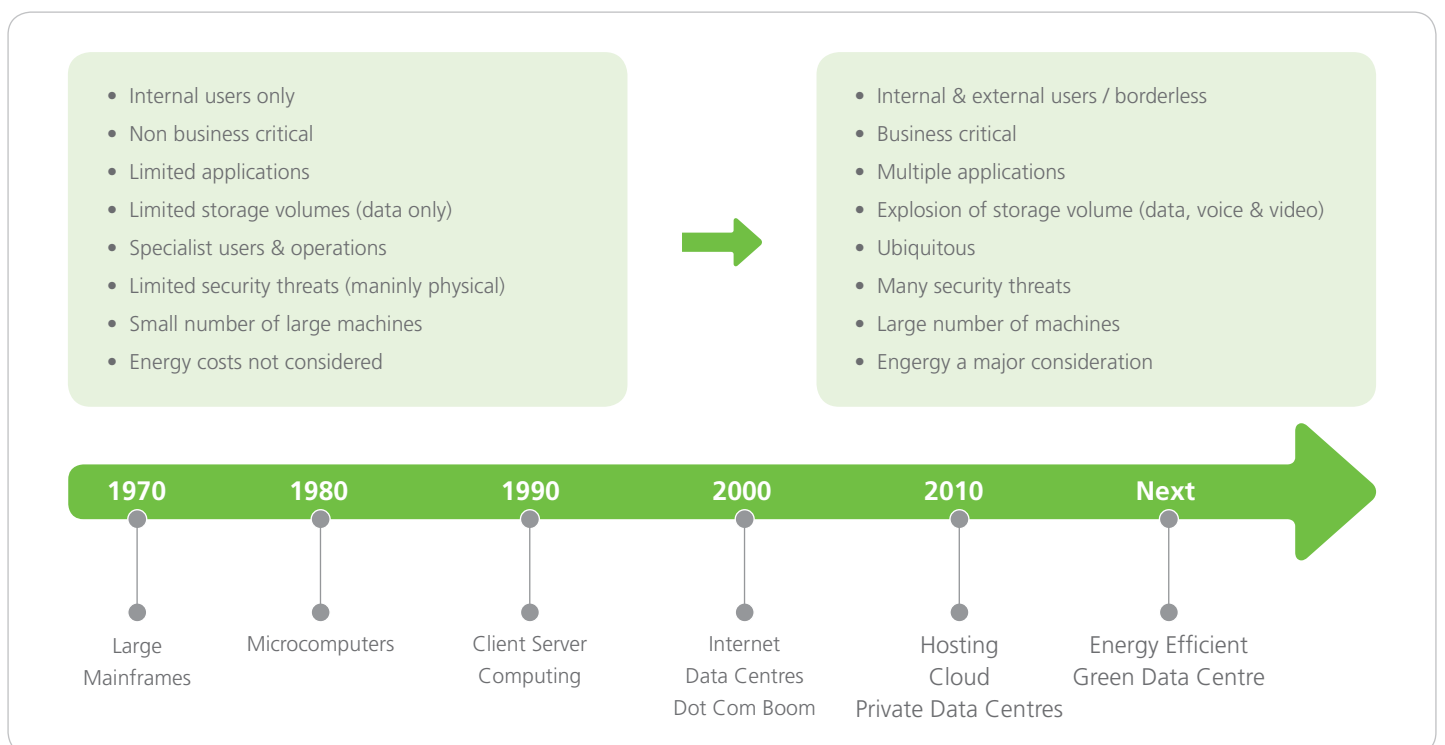
In the relatively short time span that IT has been deployed in business we have seen massive changes in the role of the data centre and the demands put upon it. In a 40-year period, the environment has changed massively, as figure 1 below highlights.

But that's not the end of the story; organisations continue to experience rapid IT growth, while the data centres required to support this growth are aging. In May 2011, data centre research organisation Uptime Institute reported that 36% of the large companies it surveyed expect to exhaust IT capacity within the next 18 months.

"Aging data centres and infrastructure have considerable energy requirements and unnecessary operational costs, imposing capacity limits that affect uptime and flexibility. Poorly managed real estate and underutilised technology also lead to inefficiencies and high energy costs."

– Helen Tang Worldwide Lead, Data Centre Transformation Solutions HP Enterprise Business

Figure 1: IT Evolution and the data centre



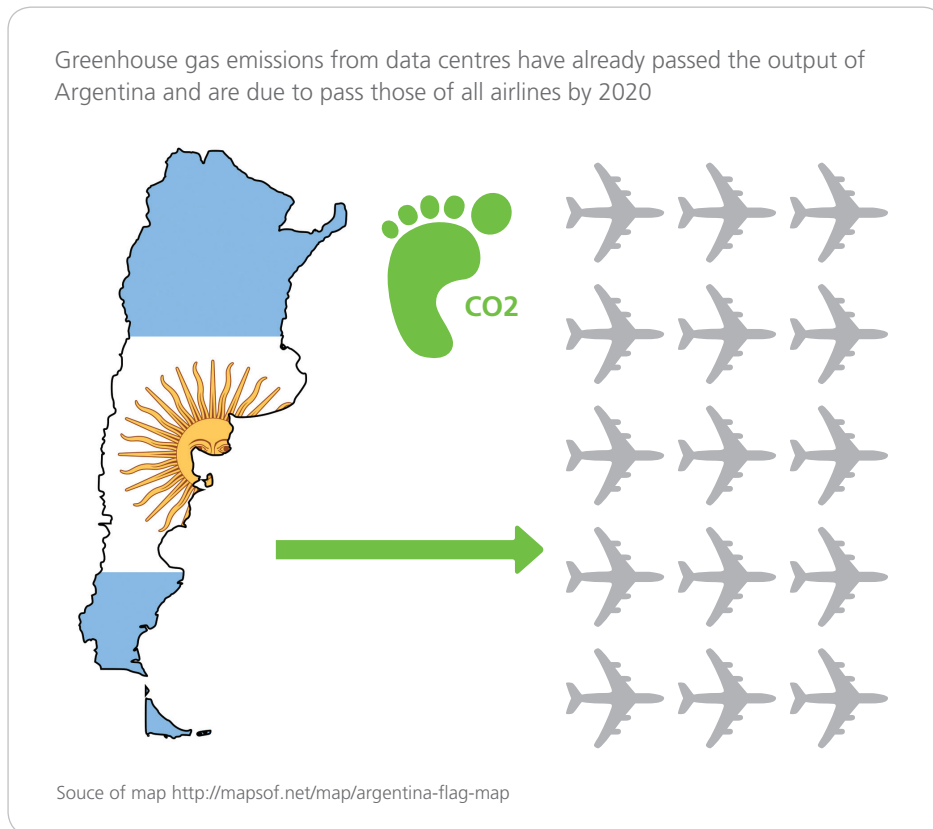
Hence, the stage is set for the need to build more data centres to cope with the capacity demand, and/or upgrade and expand existing data centres, and – in either case – ensure that the current energy inefficiencies are addressed. The facts and figures are:

- A Stanford University report published in 2011, entitled 'Growth in Data Center Electricity Use 2005 to 2010', estimated that the amount of electricity consumed by data centres worldwide grew by 56% between 2005 and 2010.

- IT energy costs are increasing by 16% every year, according to a report by McKinsey & Company and the Uptime Institute. Electricity used in global data centres in 2010 is likely to have accounted for 1.1 – 1.5% of total electricity use, respectively. For the US, that number was 1.7 – 2.2%, according to the Stanford study.

Unless addressed, this will also influence carbon emissions – to which IT data centres are a major contributor – and, where applicable, the ability of a business to meet regulations and commitments.

Figure 2: Data centre carbon footprint



IT energy costs are **increasing by 16% every year**, according to a report by McKinsey & Company and the Uptime Institute. Electricity used in global data centres in 2010 is likely to have accounted for 1.1 – 1.5% of total electricity use, respectively.

What can be done?

The good news is that there are ways to address this excessive use of energy, but identifying the optimum combination of approaches to address the issue in your organisation is needed. The answers are typically split between technology infrastructure approaches and facilities infrastructure approaches. Technology infrastructure is often offered by manufacturers and resellers; facility infrastructure is generally offered by specialist facility builders and infrastructure equipment vendors, with facility services usually offered by specialist providers.

Since suppliers tend to function in one of these two areas, only half of the solutions are presented for consideration.

McKinsey created a 10-point list of areas to address in their report entitled '10 Game Changing Improvements to Double Efficiency'.

Figure 3: '10 Game Changing Improvements to Double Efficiency'

Facilities infrastructure

- Create data centre energy dashboard to identify pain points
- Reduce cooling loss through sealing cable cut-outs, rationalising cable plans, etc.
- Increase cold aisle temp, eliminate 'duelling' cooling units, increase chilled water temp
- Implement free cooling where applicable
- Selectively turn off core components to increase remaining unit efficiency

Technology infrastructure

- Remove dead servers
- Upgrade older equipment (non-virtualised)
- Deploy virtualisation for existing and new demand
- Reduce demand for new servers
- Introduce greener and more power efficient servers and enable power-saving feature

Source: McKinsey and Company

For example, an average data centre runs at 18-21°C. However, the majority of hardware can easily run at 30°C. Therefore, as Quocira noted in its report entitled 'The Future of Green – February 2010', running at 26°C will have minimal impact on the equipment's life, but a massive impact on energy used for cooling and raised output heat quality.

One key point to emphasise about the hardware infrastructure is the impact of moving from current client server infrastructure to virtualisation, and then to unified infrastructure. The figures below show the significant reduction in server population, bringing with it the economic benefits of lower power, cooling and simpler management.

The McKinsey list provides an excellent basis upon which to create additional ways of addressing the situation.

Figure 4: Technology infrastructure Impact

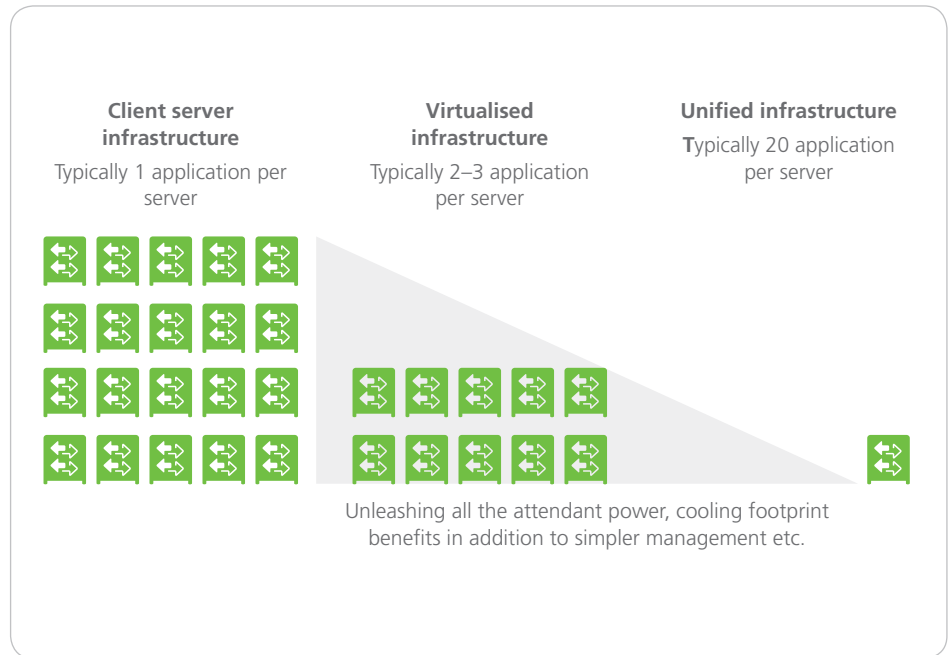


Figure 5: What can we do?



Below are definitions of the terms used in the 'Facility services' section of Figure 5, as these may be less familiar:

Cogeneration	Cogeneration – also known as combined heat and power (CHP) – is the use of a heat engine or a power station to simultaneously generate both electricity and useful heat.
Trigeneration	Trigeneration – also known as combined cooling, heat and power (CCHP) – refers to the simultaneous generation of electricity, useful heating and useful cooling from the same original heat source such as fuel or solar energy. Waste heat insufficient in energy to produce electricity through mechanical means, such as through a steam turbine, contains usable energy that's harnessed for thermal regulation purposes as described in cogeneration. Trigeneration differs from cogeneration in that some of the waste heat is used for cooling. Trigeneration systems can attain higher efficiencies per unit fuel than cogeneration or traditional power plants.
Alternative energy	Alternative energy is an umbrella term that refers to any source of usable energy intended to replace fossil fuel sources and reduce the CO ² output of the data centre without the undesired consequences of the replaced fuels. Such sources include air, water, solar, geothermal or a combination of these sources to increase the overall efficiency of the systems.
Low energy cooling	By using free cooling methods and intelligent building controls, the cooling requirements can be met at lower energy costs to those of a typical existing system. To name a few examples, these could include variable speed drives, low loss cooling towers, and evaporative cooling.
Heat recovery	In instances where the treated air still has useful energy left within it, this can be recovered or re-used instead of being exhausted.
Load Shedding & Store	Load shedding is used in times of mains failure where non-essential servers are shut down, which in turn reduces the size of back-up generators and UPS sets.

Since any combination of solutions could address your situation, working with an organisation who can take the broader view is imperative.



How do we know we are making a difference?

To measure progress, meaningful metrics need to be adopted. Lower energy consumption would be a good start, but being more sophisticated is necessary to fully understand the impact. A number of measures have been put forward, which are outlined below, and an additional measure has been added for consideration.

Power Usage Effectiveness (PUE) is a metric used to determine the energy efficiency of a data centre and was created by members of the Green Grid, an industry group focused on data centre energy efficiency. PUE is calculated by dividing the amount of power entering a data centre by the power used to run the computer infrastructure. The resulting PUE figure is expressed as a ratio, with overall efficiency improving as the ratio decreases toward 1.

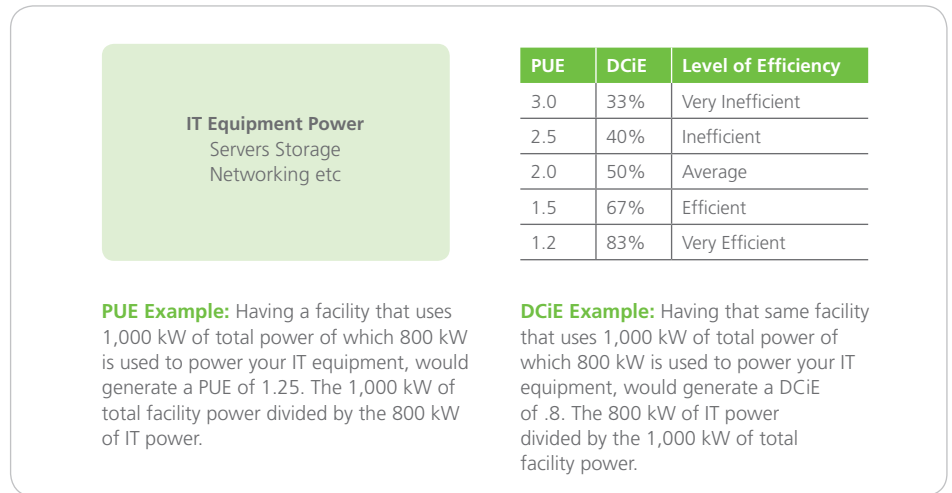
$$PUE = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

Data Centre Infrastructure Efficiency (DCiE) is the reciprocal of PUE, and is expressed as a percentage that improves as it approaches 100%. Although, the DCiE seems to be the same PUE, it uses the same numbers to express a different result.

$$DCiE = \frac{\text{IT equipment power}}{\text{total facility power}}$$

An alternative metric, proposed by McKinsey & Company and the Uptime Institute, is called **Corporate Average Data Efficiency (CADE)** and combines both IT and facilities costs to monitor energy. CADE combines measurements of the energy efficiency and utilisation of IT equipment and facilities into a single percentage. A higher CADE indicates a more energy efficient data centre.

Figure 6: PUE and DCiE examples



$$CADE = \text{Facility Efficiency (FE)} \times \text{Asset Efficiency (AE)}$$

Where:

- Facility Efficiency (FE) is equal to (facility energy efficiency) x (facility utilisation)
- Asset Efficiency (AE) is equal to (IT energy efficiency) x (IT utilisation)

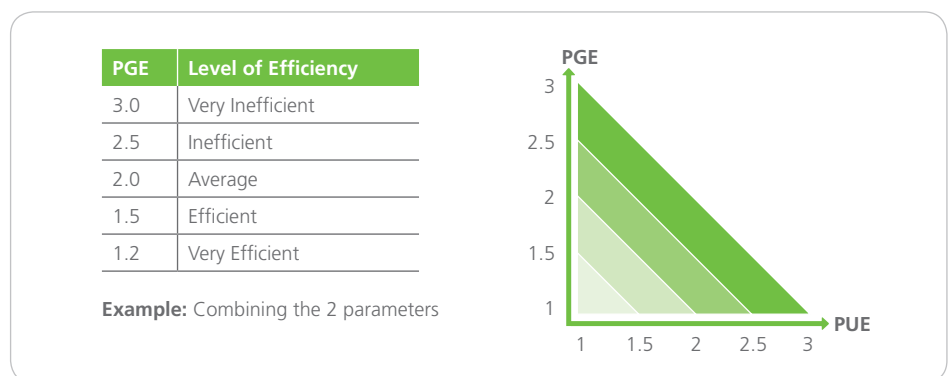
The PUE or DCiE takes into account how efficiently the energy is being used to power IT equipment, but these measurements don't calculate the sustainability of the data centre. Therefore, an additional measure could be defined, **Power Generation Efficiency (PGE)**.

$$PGE = \frac{\text{Total facility power}}{\text{total power from renewable sources}}$$

As an example, a facility that uses 1,000 kW of total power, of which 800 kW is generated from renewable energy sources, would generate a PGE of 1.25. This number is derived from the 1,000 kW of total facility power divided by the 800 kW of renewable generation.

To get an ample metric, a combination of PUE and PGE or CADE and PGE calculations would give a true figure of the overall sustainable credentials of the data centre.

Figure 7: Measuring a sustainable data centre



Lower energy consumption would be a good start, but being **more sophisticated** is necessary to fully understand the impact

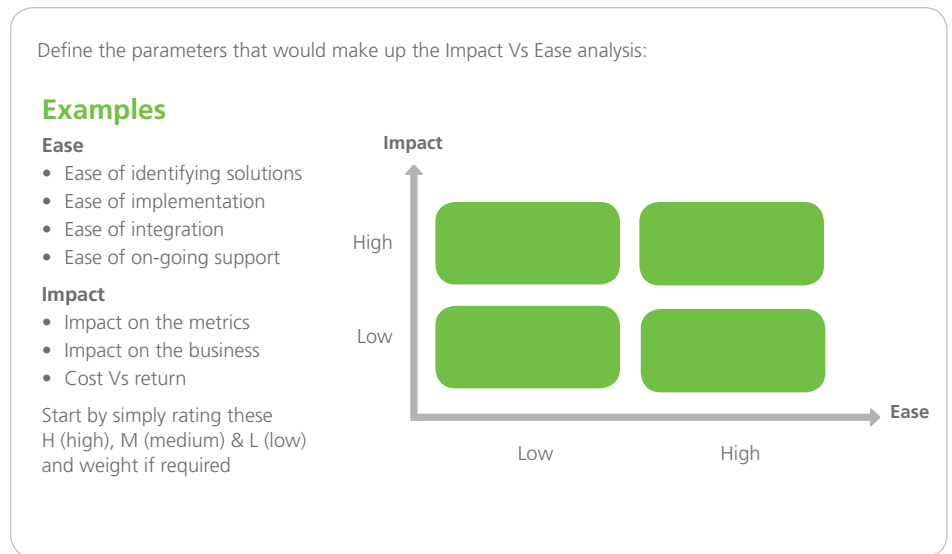
How do we get started?

Improving and measuring the sustainability of your data centre is a large and complex problem and must be tackled both from a business and environmental perspective.

Here's a recommended seven-point action plan to get started:

1. If you haven't already done so, appoint an individual with the authority, influence, credibility and respect within the organisation to take ownership and leadership of implementing solutions and monitoring your progress.
2. Establish a baseline of where you are today, using the appropriate metrics above, and set a tangible goal to achieve improvements that impact and benefit the business.
3. Set a performance indicator or measurement for the person you have appointed to deliver the improvement. Ensure that this person has the remit to work across functions within the business (e.g. IT development, IT operations, facilities, procurement, etc.).
4. Set up the appropriate governance to review progress and make sure organisational boundaries don't affect this progress.
5. Identify areas for improvement and build a plan with clear milestones – looking at the impact and ease of addressing each area of improvement.
6. Implement the plan and work with organisations that share your holistic, cross-functional approach.
7. Monitor and measure the impact on the agreed metrics and reference this against the baseline (step 2) to ensure you are achieving the desired goals.

Figure 8: Impact versus Ease analysis



In conclusion

Understanding and reducing energy in an organisation's data centre is a very real problem that must be addressed. This will influence the ability of a business to deliver cost-effective IT and meet regulations and commitment targets. It can be addressed by understanding what's available, how to measure the results, and how to progress from there.

Understanding and reducing energy in an organisation's data centre is a very real problem that must be addressed.

Figure 9: Call to action

To Do ...
<p>Drive Down IT Costs</p> <ul style="list-style-type: none"> • Drive down energy costs • Every server virtualised saves up to 7000kWH of electricity annually
<p>Increase Asset Utilisation</p> <ul style="list-style-type: none"> • Improve server utilisation Running fewer, highly utilised servers frees up space and power • Reduce your number of servers eliminating server sprawl and cutting maintenance costs
<p>Reduce Carbon Emissions</p> <ul style="list-style-type: none"> • For every kWH eliminated saves 0.5kg CO²

MIDDLE EAST & AFRICA

ALGERIA • ANGOLA
BOTSWANA • CONGO • BURUNDI
DEMOCRATIC REPUBLIC OF THE CONGO
GABON • GHANA • KENYA
MALAWI • MAURITIUS • MOROCCO
MOZAMBIQUE • NAMIBIA • NIGERIA
RWANDA • SAUDI ARABIA
SOUTH AFRICA
TANZANIA • UGANDA
UNITED ARAB EMIRATES • ZAMBIA

ASIA

CHINA • HONG KONG
INDIA • INDONESIA • JAPAN
KOREA • MALAYSIA
NEW ZEALAND • PHILIPPINES
SINGAPORE • TAIWAN
THAILAND • VIETNAM

AUSTRALIA

AUSTRALIAN CAPITAL TERRITORY
NEW SOUTH WALES • QUEENSLAND
SOUTH AUSTRALIA • VICTORIA
WESTERN AUSTRALIA

EUROPE

BELGIUM • CZECH REPUBLIC
FRANCE • GERMANY
ITALY • LUXEMBOURG
NETHERLANDS • SPAIN
SWITZERLAND • UNITED KINGDOM

AMERICAS

BRAZIL • CANADA • CHILE
MEXICO • UNITED STATES