Green Sustainable Data Centres

Measurement and Control

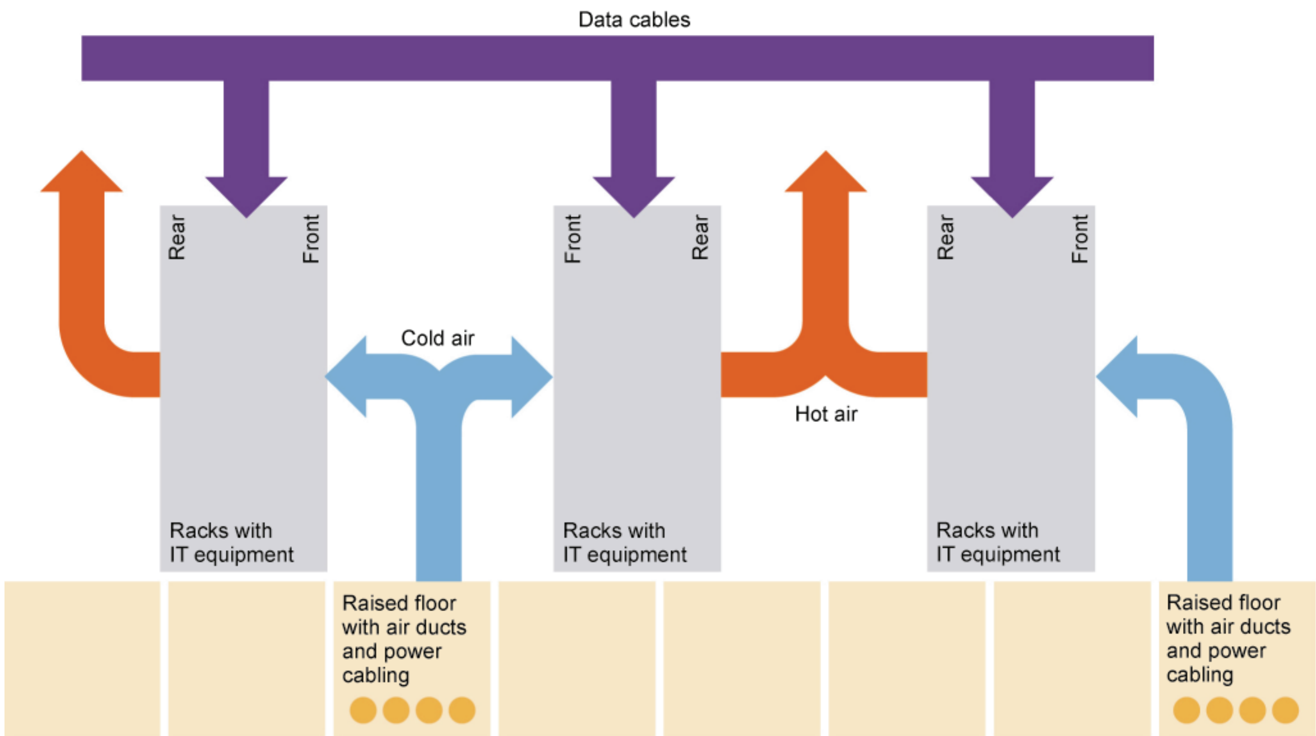

This course is produced under the authority of e-Infranet: http://e-infranet.eu/

Course team
prof. dr. Colin Pattinson, Leeds Beckett University (United Kingdom),
course chairman and author of Chapter 1 and 7
prof. dr. Ilmars Slaidins, Riga Technical University (Latvia*),*
assessment material development: Study Guide
dr. Anda Counotte, Open Universiteit (The Netherlands),
distance learning material development, editor- in-chief
dr. Paulo Carreira, IST, Universidade de Lisboa(Portugal),
author of Chapter 8
Damian Dalton, MSc, University College Dublin (Ireland),
author of Chapter 5 and 6
Johan De Gelas, MSc, University College of West Flanders (Belgium),
author of Chapter 3 and 4
dr. César Gómez-Martin, CénitS - Supercomputing Center and
University of Extremadura (Spain),
author of Checklist Data Centre Audit
Joona Tolonen, MSc, Kajaani University of Applied Sciences (Finland),
author of Chapter 2

Program direction
prof. dr. Colin Pattinson, Leeds Beckett University (United Kingdom),
prof. dr. Ilmars Slaidins, Riga Technical University (Latvia)
dr. Anda Counotte, Open Universiteit (The Netherlands)

Hosting and Lay-out
http://portal.ou.nl/web/green-sustainable-
data-centres
Arnold van der Leer, MSc
Maria Wienbröker-Kampermann
Open Universiteit in the Netherlands

First edition 2014

Measurement and Control

Chapter 5

Measurement and Control

Damian Dalton
University College Dublin

INTRODUCTION

In Chapter 4 we discussed energy savings by consolidation of servers. The next step is consolidation of data centres and to use software and hardware of the cloud. Many data centres are arranged to serve this need. A second tendency is green procurement, data centres or cloud services is an area where these criteria can be applied.
In this chapter we will first discuss the evolution of the cloud and then we will zoom in to the real subject of this chapter: measurement and control: what and how to measure in a data centre to control energy efficient operations. We will argue why a DCIM system is very attractive in order to control the data centre. The details of a DCIM system will be discussed in Chapter 6.

LEARNING OBJECTIVES
After you have studied this chapter we expect that you are able to:
– know the definition of standard metrics and understand the purpose on their use in data centre management
– know the main energy performance indicators and interpret energy data reports and dashboards
– calculate these metrics using conventional tools or methods
– use these metrics to increase performance and efficiencies in the data centre
– know how to incorporate metrics into control mechanisms regulating cooling and power distribution in order to reduce operating costs.
– use metrics in developing and driving a green, sustainable policy in a data centre
– understand the current limitations with existing metrics and the directions and actions that various standards' bodies are taking to address them.

Study hints
The purpose of this chapter is to discuss the role and importance of proper and effective energy management in the data centre and the metrics which are used for this purpose. Papillon, introduced and demonstrated in this chapter affords the opportunity to use and explore a state-of-the-art energy monitoring and management system on a live rack/server system, and learn how to extract and interpret energy-related data and information with the intention of reducing costs and increasing energy efficiencies. While reading this chapter the student should reflect on the flowing themes and issues:
– Apart from costs why is energy and power monitored?
– For each metric that is introduced, what is its function or purpose?

– What data is required to produce each metric?
– Are all metrics of equal value or importance?
– What energy-savings actions, if any, are highlighted or suggested by each metric or measurement?
The workload is 12 hours.

CORE OF STUDY

1 Cloud Computing

The cloud has evolved because of advances and developments in three principle areas:
1 A cheap and effective, global communication system—The Internet.
2 Memory and Storage devices now have enormous capacities and are inexpensive.
3 Computer processing power has doubled every 18 months since the 1960's (Moore's Law) and costs have reduced at an even faster rate.

Big Data

These developments have precipitated an avalanche in the number and variety of devices that are now connected to the internet, with the result that in addition to the 2.4 Billion humans that use the internet with some kind of computer (PC, laptop, tablet, smartphone), there are 100's millions of other computer devices generating data and information every day. This is called 'The Internet of Things'. By 2020, it is anticipated that 50 Billion (non-human) devices will be communicating on the internet. Already, a massive daily injection of information is deposited into the cloud on everything from what people are eating for their breakfast, to tracking the spread of epidemics and diseases by tracing the Google logs of a country's population. This is the Era of '*Big Data*' and a whole new industry is evolving to provide the knowledge and technology to mine these deposits and extract whatever information can be gathered using *correlation* techniques, rather than the conventional engineering approach applying the logic of *causality*. This is the Science of Data Analytics.

Computing engines

Data centres are the *computing engines* of the cloud; large, anonymous warehouses filled with avenues of racks. Each rack holds between 10 to 20 servers and centres have anything from 100's to 10,000's of racks. There are approximately 500,000 data centres world-wide.

These new, modern wonders of the world come with an environmental cost. They are vast consumers of electricity with the same carbon footprint as the global airline industry. On average, for every $2 spent on powering a server $1 is spent keeping it cool. It is not surprising that northern latitudes, where cool air comes free, have a climatic advantage over certain other more southerly locations. Furthermore, it is politically desirable and frequently cheaper to source green energy such as hydroelectricity to power data centres. These are some of the reasons why Facebook located its new data centre, the size of 11 football pitches at Lulea, Sweden. For numerous reasons political, financial or logistical, the location of data centres is not determined solely by the coolest locations with the greenest or cheapest energy. Data centres are ubiquitous and making them as energy efficient as possible is a common objective in all decisions for all locations.

Overall, the cloud is a new computing methodology, offering exciting new commercial and business solutions and opportunities, but in doing so it introduces other challenges particularly in the energy domain that must be addressed.

1.1 INTERESTING FACTS CONCERNING THE CLOUD [1]

– Google processes more than 24 petabytes (2^{50} bytes = 10^{15} bytes approx) per day, a volume that is a thousand times the quantity of all printed material in the U.S Library of Congress.
– The 800 million users of YouTube upload over 1 hour of video per second.
– Facebook members 'Click' or comment 3 billion times per day.
– In 2007, 300 exabytes (2^{60} bytes = 10^{18} bytes approx) of stored data existed world-wide. This is doubling every 3 years.
– The storage capacity needed by the average Fortune 1000 company doubles every 10 months

The deficiencies of conventional power monitoring technologies are contributing to data centre power demands spiralling out of control. The enormity of this power problem is highlighted in the following statistics and actions [2]:
– Typically for every $2 spent on server power, $1 is spent on cooling it.
– In 2005, 1.5% of the U.S electricity was consumed by server farms in data centres. This amounted to $4.5 billion worth of electricity, or roughly 61 billion kW hours, the equivalent of 5.8 million US households. In 2011, the total U.S data centre energy bill was $7.4 Billion.
– In 2005, $26.1 billion was spent powering and cooling the global installed server base. Over the next 5 years this grew at 11.2% CAGR.
– $41.4 Billion in global revenue (28% of the total data centre market) will be spent on the Green agenda in data centres over the next 5 years.
– Data centre energy consumption worldwide has doubled since 2000. There are now 35 million servers worldwide.
– The electricity consumption in 2011 of all European data centres was equivalent to that of Portugal and is expected to double by 2020.
– By 2020, it is predicted that the Carbon footprint of the EU data centre community will constitute 15-20% of Europe's total CO_2 emissions.
– In the EU Code of Conduct on Data Centres Energy Efficiency (ISPRA 30/10/2008) the EU-countries agreed on a single EU-wide cap on emission allowances which will apply from 2013 and will be cut annually, reducing the number of allowances available to businesses to 21% below the 2005 level in 2020.

1.2 CLOUD COMPUTING DEPLOYMENT MODELS[1]

Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

[1] NIST. (2011). NIST Definition of Cloud Computing.
http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf

Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
Public cloud. The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds). The cloud infrastructure supports a specific community consisting of several organisations that have shared interests (mission, security, policy). Management is carried out by the organisations themselves or is outsourced to a professional IT service provider.

1.3 POWERING THE CLOUD

Reliability
Capacity
Delivery of service and performance

Historically, data centre design and operations have been focused on *reliability* and *capacity*. This has led to the unfortunate situation where data centres have not been optimized for energy efficiency. The main focus of attention for the directors, managers and operations staff, and those charged with the design and operation of a centre has predominantly been *delivery of service and performance*. Until fairly recently, the task of energy efficiency has been an intention but not a single responsibility for any of the main stakeholders in the data centre organisation. Only as costs have escalated has energy efficiency now become a priority.

Necessary power
Downtime

It is not simply a question of reducing costs and the climatic impact of data centres the national power grid of some European countries can no longer supply the *necessary power* to the computer systems in banks and large organisations in many major cities. For practical purposes, many new computing resources must be located outside urban areas. Despite the awareness of the chronic threat of an electricity supply shortfall, most senior data management were found in a 2010 survey to be oblivious to the real vulnerability of their data centre. The Ponemon Institute study [3] of 400 U.S data centres, commissioned by Emerson Network Power highlighted the misconceptions about the frequency and impact of data centre *downtime*. It found that the typical U.S data centre on average had 2 downtime events over a two year period due mainly to power and cooling problems, and that the average cost to the centre was $505,000 and the recovery time was 134 minutes. Ignorant to these facts, 62% of senior management believed such events only happened rarely. This illustrates the potential grave economic consequences to businesses posed by the growing power demand from ICT devices.

In a more recent 2012 survey of 2000 data centres worldwide[4], the Uptime Institute, the U.S largest professional data centre organisation, concluded with the following assessments:

1 32% of data centre 2012 budgets are at least 10% greater than 2011. Data centres are still spending despite the global recession.

2 30% of data centre facilities will run out of power, cooling and/or space in 2012. This has been a recurrent trend.

3 Reducing data centre energy consumption is very important in 71% of data centres.

4 The top driver for pursuing energy efficiency is financial.

5 The average PUE (Power Utilisation Efficiency) metric, a measure of the efficient use in data centres is 1.8 to 1.89. This implies that energy overheads are still very high, 90% relative to the direct energy consumed. There is plenty of scope for energy reductions.

1.4 POTENTIAL ENERGY AND COST SAVINGS WITH ENERGY MANAGEMENT

On a positive note, a 2008 report by McKinsey [5] has shown that enforced corporate energy management policies can reduce power consumption by up to 40% (see Figure 1). The degree of energy saving depends on the stringency of enforcement and extent of the organisation's energy saving policy: in figure 1 these are divided in:
– Monitoring and measurement
– Retrofit and repair
– Control demand
– Energy performance management.

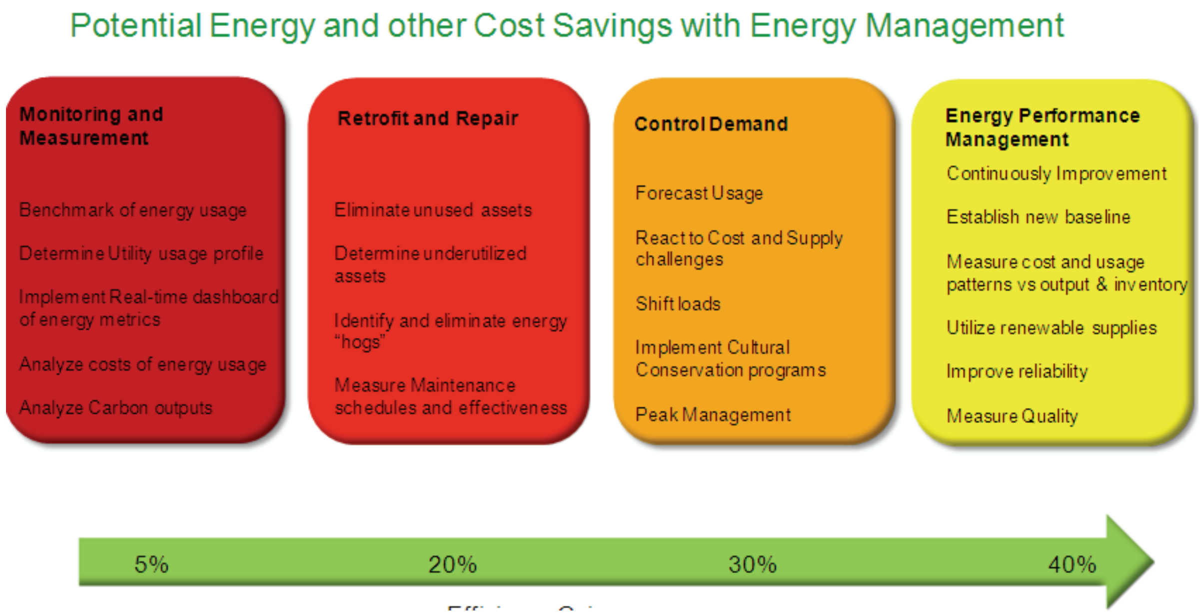

FIGURE 1 Significant reduction (increased efficiency) in data centre energy consumption is possible when the proper energy monitoring tools enforce an informed corporate energy policy

The evidence in all the reports concerning the location, operation and management of data centres around the world, indicates the imperative to provision adequate and secure energy supplies for current and future demands, and the necessity to monitor them with integrated management systems termed DCIM (Data Centre Infrastructure Management) tools (which will be discussed later). These tools ensure that data centres are reliable, efficient and resilient to internal and external adverse conditions and circumstances. Amongst the tasks regulated by these tools, most commentators would single out power management as the top priority in any data centre agenda, as any inadequacies may jeopardise a centre's financial viability or its operational integrity.

2 Motivation for Energy Management in the Data Centre [6],[7]

2.1 REMOVAL OF OVERSIZING

Oversizing, where the manufacturer's 'Name-plate' energy rating is used in provisioning adequate PDU power, leads to one of the largest operational inefficiencies in the data centre. Typically, the manufacturer deliberately over-estimates the power requirements of their servers, so that the data centre makes sufficient allowance in their power allocation. Unfortunately, the manufacturer ratings can be 50% or more than the server's actual maximum power consumption. See Figure 2. The result is that the power infrastructure is over-provisioned, particularly the PDUs. The servers, collectively in each rack consume a lot less power than expected, and consequently the PDUs operate at a much lower power demand where they are less efficient (70-80% instead of < 95%). Furthermore, the racks are not as tightly populated with servers as possible since the assumed power demand is much higher than in reality. The net effect is that capital expenditure is wasted on larger PDUs than necessary, the PDUs operate more inefficiently, the capacity of the PDUs is never fully utilised, and more rack and floor-space is occupied than really required. The name-plate information should be adjusted to reflect the real power consumption of the server or any device or equipment in the data centre.

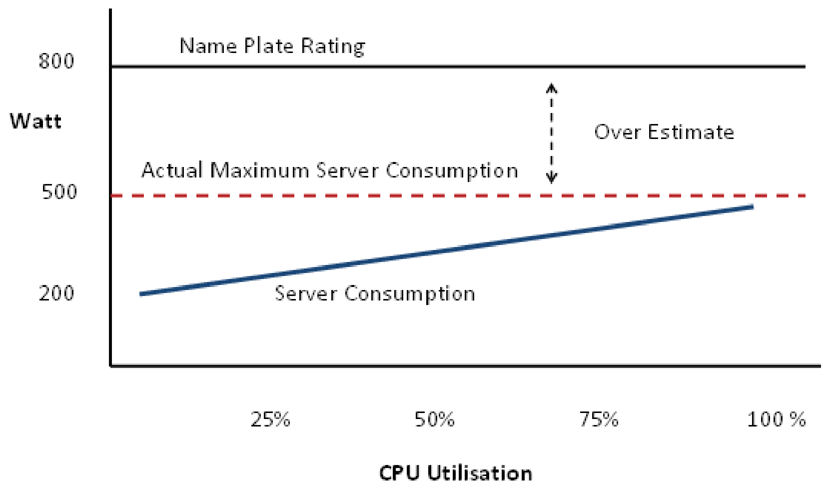

FIGURE 2 De-rating Equipment is Essential to Avoid Over-Provisioning for Power

Right-sizing

Right-sizing the energy demands of a data centre can save up to 30% of energy costs and substantially reduce the cost of actual real-estate that is provisioned. The cost of building a data centre (2011) per m^2 of data centre space in the U.S ranges between $5,000 (Tier 1) to in excess of $13,000 (Tier 4), this figure can be used to estimate the savings by reducing rack space. For the Tier classification see Chapter 2.

Stranded Power

Average
Actual Peak power

Maximum possible power consumption

Safety margin

Stranded

A similar over-estimate in rack power provisioning, again usually as a consequence of over-estimation of the rack's indigenous servers or a mismatch between the power capacity of PDU and the power demand of a rack, leads to the concept of *Stranded Power*. Several power observations can be made of a rack while its servers execute their tasks. From monitoring the actual collective power activity of the servers, the *Average* and *Actual Peak power* consumption of the rack can be measured. The observed actual peak power consumption is usually 20-50% less than the *Maximum possible power consumption* than all the servers can actually draw. However, when power for racks is being specified, a worst case analysis is usually adopted to give a *safety margin* of operation, and in determining power requirements the maximum possible power consumption of each server is assumed plus further margin of about 20%. This produces a specification which is grossly over-estimated demand which is never realised, the difference between what is provisioned and the maximum possible power consumption is available but not recognised, this is *Stranded* power and can account for approximately 75% of a rack's power capacity. See Figure 3. Stranded capacity is a similar phenomenon which can also apply to floor and cooling capacity and originates in many cases from inadequate power knowledge.

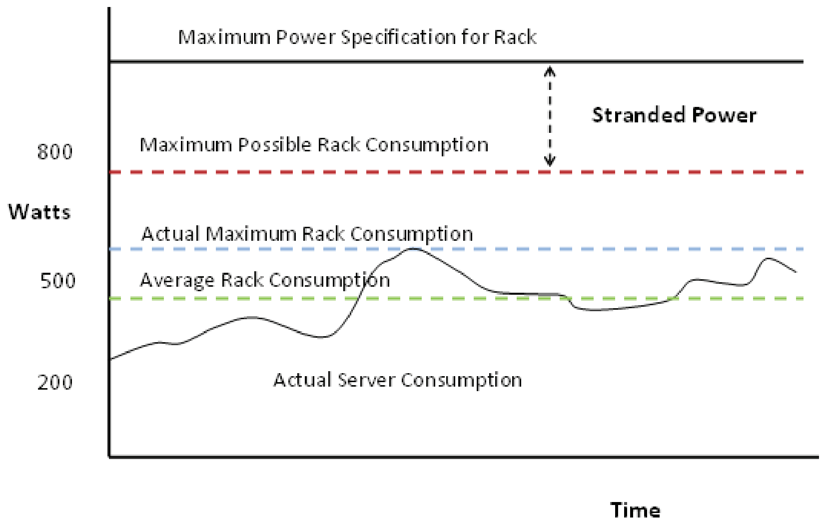

FIGURE 3 Over-estimation of Power Requirements Leads to Stranded Power

2.2 IMPROVING EFFICIENCY

In a typical data centre maintaining the correct electrical, environmental and climatic conditions for the sustained and satisfactory operation of the IT equipment is a major logistical and physical overhead. See Chapter 2 for a discussion of the configuration of a data centre.

Uninterruptable Power Supply (UPS)
Power Distribution Unit (PDU)

Overhead

Electrically, power from the external grid or major power generator is distributed into the data centre by a *UPS* (*Uninterruptable Power Supply* and then to the racks via a *PDU* (*Power Distribution Unit*). Environmentally, the temperature and humidity is regulated by Chillers, Humidifiers and the CRAC (Computer Room Air Conditioning)/ HVAC (Heating Ventilation Air Conditioning) systems in the data centre. Collectively, these are *overheads* which for practical and economic reasons any data centre will attempt to reduce as far as possible. The proportion of the total data centre energy consumed by the infrastructure reduces as the IT becomes more loaded, in other words the centre becomes more efficient. See Figure 4. Metrics defining efficiency and the degree of inefficiencies in data centres will be discussed later.

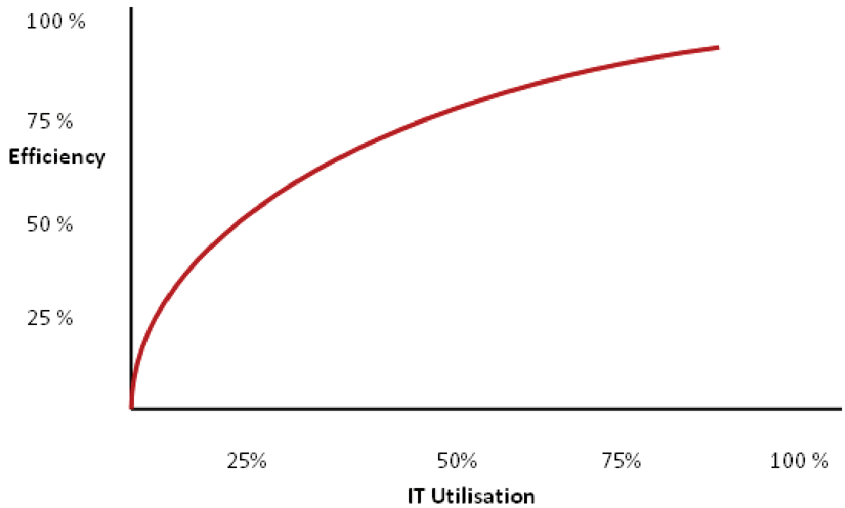

FIGURE 4 Data Centre Efficiency Improves with IT Equipment Utilisation

An example of how efficiency improves with increasing IT loading is the typical power performance behaviour of servers. An idle server can consume 50% of its peak power. Moving from 10% to 50% utilisation, a 5-fold increase, only expends 40% more power.

Task 1
Review Figures 1 and 5 in Chapter 2

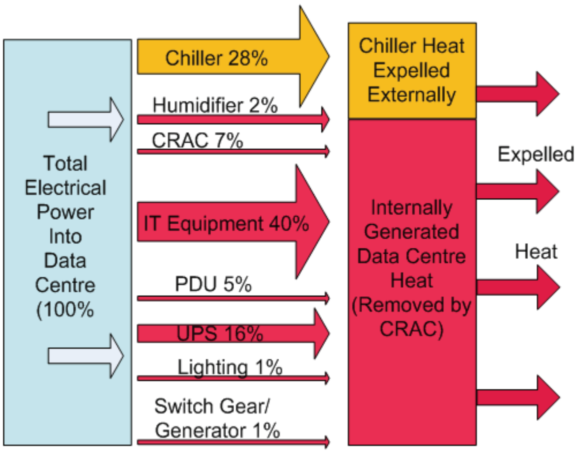

FIGURE 5 Heat Generation in a Typical Data Centre

Power -> heat
Figure 5 illustrates power dissipation in a typical data centre. The electrical power from the grid, or in some cases the centre's own generator, is distributed to the IT equipment by the UPS and PDUs. Virtually all power in a centre, including the IT equipment, is dissipated as heat, so the power inefficiencies of the UPS and PDUs, 16% and 5% respectively will manifest it-self as internal heat in the data centre. The IT equipment, operating at a typical 20-40% average utilisation, will consume in the region of 40% of the total energy delivered to the data centre. This is the only productive use of all the power supplied, the remaining 60% is the overhead power whose sole purpose is to provide a working environment for the IT equipment. 72% of the power entering the centre will generate internal heat which must be removed, the CRAC and humidifier will perform this task but they themselves will contribute some internal heat in the process. The internal heat is removed from the server racks and all other equipment by the CRAC system which relies on a Chiller unit to cool the ambient external air before it is blown through the centre and then expelled with the centre's heat to the exterior. The Chiller does not contribute to the internal heat of the centre as its heat is dissipated externally by water, however, major energy savings are made with efficient Chiller and CRAC systems, or in geographical situations where the ambient temperature of the external air is so low that it is not necessary to chill it, it is only required to push the air through the centre.

Real-time monitoring

Not feasible to monitor every device

Since the core activity of the centre is data processing, in an ideal situation, this activity would consume all the power supplied. However, while this is a target that can never physically be achieved, nonetheless, it is in the interest of every centre to constantly strive to reduce operating costs and move towards this ideal. Realistically, this reduction can only be accomplished by adequate *real-time monitoring* of the power consumption of all equipment and devices in the centre, so that minimal but sufficient power and cooling can be allocated to where it is needed, when it is needed. While the real-time information must be accurate, at the same time it is *not feasible*, or even necessary, *to monitor every device*. This topic will be reviewed later.

Blades

Performance (Instructions per watt)

Higher power-density racks

Hotspots

In the context of reducing the large cooling costs, inaccurate power information leads to a loss of clarity in focussing cooling to where power and hence heat is being generated. One must rely on detecting secondary effects such as temperature increments which is less desirable and efficient. This issue is becoming more significant with advances in server technologies and as server *Blades* (see Chapter 3) become more mainstream. For example, multi-core processors based on Intel's Core micro-architecture deliver approximately 5 times more *performance (Instructions per watt)* than single core processors based on the earlier Intel NetBurst micro-architecture. Blades offer higher performance but also execute proportionally more instructions than conventional servers, so that they actually consume more power per unit. By their design, they have fewer components such as power units per board, so physically they occupy less space than many server types. A blade enclosure in a rack occupies between 4U and 8U in height but can take 8 to 16 blades. This concentration of servers has led to *higher power-density racks* which must be more closely regulated as regards cooling and power provisioning so as to avoid the generation of *Hotspots* in data centres, regions in a centre which are inadequately cooled and which can culminate in server failures. Table 1 Illustrates the power-density challenge, the results are from a Ziff Davis Enterprise, eWeek publication 2009 survey of IT companies.

TABLE 1 Power-Density Growth in Data Centres

Year	*Av. Watt/ft^2*	*Av. kW/Rack*	*MegaW consumed*	*Annual Utility Cost*
2003	40	2	4	$288,000
2005	80	4	8	$576,000
2007	240	15	24	$1,700,000
2009	500	30	50	$3,600,000

In just 6 years the average power requirement per rack has grown 15-fold from 40 W/ft^2 to 500 W/ft^2. This is likely to grow even more substantially over the next few years as servers become more virtualised and the utilisation rate of the servers increases.

2.3 ENERGY MANAGEMENT POLICY

Accurate power and energy monitoring is also an integral part of any energy management policy in an organisation. The maxim "If you don't measure it, you can't manage it" applies. Measuring energy at various levels in a data centre is the only technique which can enforce and validate the merits and advantages of any green, sustainable or even a simple energy reduction policy as it is delivered through a number of strategies. Many reports confirm the observation that 30% of servers have a utilisation rate of 3% or less, even at this level of utilisation servers can consume 50-70% of their maximum consumption. However, this realisation is often missed and unaccounted in operational audits. Putting a financial figure on this inefficiency by actually measuring the operational energy costs is often essential to trigger any action to reduce it.

REFLECTION 1

a What are the motives for energy management in a data centre?
b What are the issues?

3 Energy Management and its Relationship to DCIM

DCIM combines information from a multitude of sources into a single point of contact. All details on the data centre physical infrastructure, the assets, their location and use, together with the centre's operational conditions, energy use and climatic environment is assembled and processed within the DCIM platform. A number of energy efficiency metrics have been standardised by international standards bodies like the Green Grid and ASHRAE such as PUE (Power Usage Effectiveness), its reciprocal Date Centre Infrastructure Efficiency (DCiE) and Air Conditioning Airflow Efficiency (ACAE)[2] which are used as indicators to measure any improvements in the power or cooling behaviour of the centre. There are several metrics commonly used, none by themselves can adequately convey the complete energy or efficiency state, but several taken in conjunction with others can give a fairly good overview of the state of the centre from different perspectives. As improvements are made in the centre, their effect can be assessed by these metrics. The data involved in calculating these metrics requires typically periodic sampling of numerous physical parameters such as electricity consumption, temperature, humidity, airflow from the IT equipment and infrastructure facilities controlling cooling. Since access to this data is provided by the centre's DCIM system, it is apparent that the precision and validity of any metric is heavily dependent upon the extent of deployment and comprehensiveness of the DCIM deployed. A DCIM system that has limited visibility of a centre's activity will generate metrics of limited value.

3.1 DRIVING FACTORS BEHIND ENERGY EFFICIENCY

As we have seen in Chapter 1 prevention of climate change is the primary driving factor behind energy efficiency, however other factors may have a more immediate impact.

Increasing energy efficiency in an organisation has obvious advantages, lower energy costs, but there are other benefits as mentioned previously which are secondary such as rack consolidation and reduction in data centre real-estate. These are positive incentives which lead directly to cost savings on an organisation balance sheet. However, there are, or will be penalties imposed on organisations which do not reduce the carbon footprint of their operations. National legislation or European directives, ultimately originating from the international 1997 Kyoto agreement, will enforce a reduction in $CO2$ emissions.

In Chapter 7 we will give a systematic overview of regulations and legislations.

[2] The amount of heat removed per unit (by volume) of airflow.

3.1.1 *European Initiatives*

An Energy Policy for Europe specified a target of saving 20% of the European Union's energy consumption compared to projections for 2020. This has formed a key ingredient in the EU Energy and Climate Change Package agreed at the European Council in December 2008 (i.e. 20% efficiency improvement, 20% renewable energy penetration and 20% greenhouse gas emissions reduction by 2020). This target is not currently binding and a method for calculating the national targets has not been finalised by the European Commission (EC).

The EU Code of Conduct on Data Centres has a voluntary 5-point plan [EU Code of Conduct on Data Centres Energy Efficiency (ISPRA 30/10/2008)]. For data centres to comply they must observe the following objectives and actions.

1 Put power metrics (PUE, DCiE etc) in place.
2 Set practical targets.
3 Monitor and manage energy use.
4 Apply energy efficient technologies.
5 Measure effects (Carbon Credits).

While there is no direct financial inducement for companies or organisations to adopt this 5-point plan, being a member does confer extra social corporate credibility in the public domain and usually in the process of compliance cost-efficient energy savings are a consequence.

Green Public Procurement (GPP)

Many European public authorities now incorporate *Green Public Procurement (GPP)* (e.europa.eu/environment/gpp) criteria in their tendering processes for services and goods. The initiative is still in its infancy but selection criteria is being developed that will be used for different services and goods in determining their environmental impact through their entire lifecycle. Data centre or cloud services is an area where this criteria can be applied, and centres that are demonstrably greener or which have a more sustainable business model will be given preference.

3.2 DATA CENTRE METRICS FOR POWER EFFICIENCY [8],[9]

Power Usage Effectiveness (PUE)
Data Centre Infrastructure Effectiveness (DciE)

The most commonly quoted energy parameter expressing a data centre's overall energy efficiency is the *PUE (Power Usage Effectiveness)* and its reciprocal *DCiE (Data Centre Infrastructure Effectiveness)*.

In Figure 5 of Chapter 2 we have seen where these metrics have to be sampled in the data centre.

REFLECTION 2
What are the values of PUE and DCiE of figure 5?

3.3 IMPROVING PUE/DCiE RATINGS

In Chapters 2 and 3 we have seen the separate techniques to improve the energy efficiency of cooling and IT equipment. Now we integrate this knowledge to improve efficiency of the data centre.

Since the PUE is the ratio between Total Facility Power and IT Equipment Power, there are essentially two ways to reduce the rating:

1 *Perversely* by increasing IT power consumption while not changing the cooling configuration, this would improve the rating. But, this would suggest using more inefficient computer equipment that would consume more energy not less. Obviously this is not a recommended approach.

2 Increase the efficiency of the cooling system, thereby reducing its power consumption and making more power available to IT. This is the preferred approach.

Scenarios

As for the desired level of improvement, the EPA in the U.S has established three different *scenarios* for data centres in the U.S. (Improved Operation, Best Practice and State-of-the-Art) indicated by a change in PUE rating to between 1.6 and 1.2 (or a DCiE of 0.6 to 0.8). The benefits in reaching this target range can be profound; for example, an improvement in PUE from 2.3 to 1.3 nearly doubles the power available for IT equipment.

With the typical data centre currently (2013) having a PUE rating of between 1.5 and 2.0 (or DCiE rating of 0.5), achieving this level of improvement likely involves a range of initiatives in most organizations, including:

Eliminating inefficiencies

– *Eliminating inefficiencies*, particularly in older equipment, throughout the cooling infrastructure.

Outside air

– Making greater use of *outside air* or other outside cooling resources to minimize the load on the Computer Room A/C system and chiller plants.

Hot/cold aisle

– Adopting a *hot/cold aisle* configuration, which may involve rearranging how equipment is placed in rows and even within individual racks.

Variable cooling

– Making greater use of *variable cooling* by adjusting fan speed on the air handler, and water flow to the individual CRAC/CRAH units.

Inlet temperatures

– Increasing cold aisle server *inlet temperatures* to 80.6°F (27°C), an increase of 2°C over the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE) previous 2004 recommendations. ASHRAE's (2008) recommended operational temperature envelope for IT equipment is now 18°C-27°C and even allows a range of 10°C to 35°C. More recently (2011), it permitted certain IT equipment to operate as high 45°C. The higher operational temperatures use less electricity in the cooling process, but unfortunately may shorten the life-span of the IT equipment.

3.4 MEASURING DATA FOR USE IN PUE/DCiE

Activity of the servers

The data centre is a very dynamic entity and the amount of power consumed depends on many factors among them being the *activity of the servers*, network and memory devices, and also the climatic conditions which will dictate the amount of air conditioning and cooling required. Therefore, in ascertaining the PUE and other efficiency metrics, it is meaningless to take one or two snapshots of the centre to ascertain the data associated with the metrics. In order to get meaningful, accurate and reliable data, for any reasonable sized centre, an automated monitoring system must be employed, essentially this means a *DCIM system*.

DCIM systems

To get a representative PUE figure for the centre, the IT and facility equipment must be monitored under light and heavy loading at different times of the day and at various junctures throughout the year so as to the centre's CRAC and HVAC systems operate over all temperature and humidity conditions.

Hourly intervals

Common practice for sampling data for use in the PUE calculation is take samples at *hourly intervals* throughout the day, but taking more frequent measurement may reveal important transient issues that are lost in samples with less resolution.

3.5 DEFICIENCIES IN A SIMPLE PUE ANALYSIS [10], [11]

While the PUE/DCiE is used extensively by the data centre community in evaluating a data centre productive use of energy, it is a measurement open to abuse or misconception. The following example illustrates this point.

At a data centre it is decided to increase the utilisation of its existing IT population. Virtualisation is introduced to decrease the number of servers in use from 1000 to 500. This has immediate benefits, money is saved on the hardware, on the energy being used to power them and also on licensing, maintenance and so on.

In the simple PUE analysis, the costs of re-engineering the data centre facility tends to militate against the cooling systems and UPSs being changed. For instance, if the cooling and UPS managed to support 1000 servers, the facilities management people may choose to leave them at no cost, as they will easily serve 500.

Virtualisation paradox

If data centre had a PUE of 1.5 – for every Watt of energy going to the servers (and storage and networking equipment), a further 0.5 Watt was being used for overheads. Virtualisation has reduced the server population's energy consumption by half, but the overheads are unchanged. Consequently, the PUE has gone from 1.5 to 2 [(0.5 + 0.5)/0.5], indicating lower efficiency, though in reality the energy cost of operating the centre has reduced by 33% (Total power consumption was 1.5, now it is 1). We call this the '*virtualisation paradox*'.

3.6 OTHER EFFICIENCY METRICS

Apart from PUE and DCiE there are several other metrics for energy efficiency. In Chapter 2 we discussed the metrics CADE and CUPS, here we discuss the Effective PUE.

ePUE (Effective PUE)

To address the type of anomalies indicated in the previous section, an amended version of PUE has been suggested which has yet to become a standard metric. Nevertheless, despite an absence of official endorsement, it being used by data centre professionals particularly in virtualised situations where it takes into account the improved CPU efficiencies of these systems. The *ePUE (Effective PUE)* is defined as:

$$\mathrm{ePUE} = \frac{\text{Total Facility Power}}{\text{Utilisation rate} \times \text{IT Equipment Power}} \qquad (1)$$

To see how the ePUE is able to reflect the improved efficiencies, taking the previous example and assuming that the original 1000 servers were operating at 10% utilisation, then after virtualisation has been introduced the 500 active servers are operating at 20% utilisation. Comparing the two situations via their respective ePUE:

$$\begin{aligned} \mathrm{ePUE}_{1000} &= \frac{1+0.5}{1\times 0.1} = 15 \\ \mathrm{ePUE}_{500} &= \frac{0.5+0.5}{0.2\times 0.5} = 10 \end{aligned} \tag{2}$$

The ePUE value of 10 is a third lower reflecting the improved efficiencies. As better an indicator of efficiencies, the ePUE can still be manipulated. Executing pointless programmes on the servers purely to increase utilisation rates would reduce the ePUE and appear to improve matters without any real advantages. In practice the increased utilisation rates may ramp up power consumption and cooling, so that even this manipulation may be detected.

Energy Reuse Effectiveness (ERE)
Energy Reuse Factor (ERF)
Data Centre Compute Efficiency (DCcE)

The Green Grid has also defined two environmentally related metrics, *Energy Reuse Effectiveness (ERE)* and *Energy Reuse Factor (ERF)*. These refer to how much surplus energy is reused e.g heat for another building. Another metric it has defined is *Data Centre Compute Efficiency (DCcE)* which takes into account the proportion of IT energy that is used in productive computing work. These have yet to be standardised and used internationally.

3.7 OTHER DATA CENTRE METRICS; SUSTAINABILITY ISSUES [12]

Total cost of ownership (TCO)

Life cycle

Data centres that have a sustainability emphasis on energy consumption, carbon emissions, and water usage have more control over their decisions on growth, location, and outsourcing strategies, while still remaining competitive and fulfilling their customers' or clients' needs. With more sustainable data centres, IT organizations can better manage increased computing, network, and storage demands and at the same time lower their energy costs and hence reduce the *total cost of ownership (TCO)* of its IT base. TCO quantifies the financial impact of deploying an IT product over its *life cycle* and includes its energy cost, which is usually the biggest factor and additionally other costs like installation, licencing, maintenance and support. The future presents risks, especially when it comes to carbon taxation and water costs and rights. Organizations that proactively focus on these issues will lower their business risks, increase their potential for growth, and better manage their environmental costs. In Chapter 1 we discussed the definitions of Green and Sustainable IT.

Sustainability

Taking this approach one can define *Sustainability* as an e-infrastructure strategy such that:
a Any energy consumption should be kept as low as possible.
b Any resource should be used as fully and efficiently as possible.
In other words wastage should be minimised.
c Timely and accurate information should be produced to assess energy usage, efficiencies and resource use (wastage) in order to direct and implement improvements.

d The full environmental and social impact of activities should be considered.
e The level of IT. resource provision should be appropriate to the task being undertaken.

Five guiding principles

These *five guiding principles* will form the framework of a definition of Sustainability in this course.

Sustainable infrastructures

Extending this definition to the specifics of *sustainable infrastructures*, it implies that *their energy consumption is kept as minimal as possible in line with the available technology.*

The debate around the environmental impact of data centres is still ongoing. A complete assessment of the environmental implications of cloud versus non-cloud computing has yet to be made. Even where data centres have employed renewable energy sources, there have been arguments that this action has limited the capacity and availability of green energy to town and urban areas within the vicinity of these centres. So the full picture of the environmental effects of the data centre in local and global communities is likely to be quite complex. However, although there may not be clear conclusions, there is the recognition that their huge energy consumption and associated cooling mechanisms (which require considerable water reserves) warrant their carbon and hydro footprint to be evaluated and reduced as far as possible. Whereas some environmental actions come at a cost to a company or organisation, when it comes to carbon and hydro footprint reduction, invariably this is a saving which even allowing for any additional infrastructure requirements is cost-effective.

There are a number of carbon and hydro-based standardised metrics for assessing the environmental impact of a data centre. CUE and WUE were introduced in Chapters 1 and 2. Because they are part of the DCIM we will discuss them here.

3.8 CUE, CARBON USAGE EFFECTIVENESS [13], [14]

Carbon Usage Effectiveness (CUE)

In combination

The Green Grid have authorised (December 2010) the use of a new metric, *Carbon Usage Effectiveness (CUE)*, to address carbon emissions associated with data centres. The impact of operational carbon usage is emerging as extremely important in the design, location, and operation of current and future data centres. When used *in combination* with the power usage effectiveness (PUE) metric, data centre operators can quickly assess the sustainability of their data centres, compare the results, and determine if any energy efficiency and/or sustainability improvements need to be made. CUE is defined as:

$$\text{CUE} = \frac{\text{Total CO2 Emissions from Total Data Centre Energy}}{\text{Total IT Equipment Energy}} \quad (3)$$

This assumes that the data centre is not producing any CO_2 emissions it-self, except those produced by its consumption of energy from the national grid. The carbon footprint of a kWh is determined by external sources to the data centre, the electricity suppliers injecting their electricity into the national grid. Their source of energy can be from a broad range of generators, hydro, sun, wind, coal, gas, and nuclear and during an average day the proportion of each source and hence carbon footprint per kWh in the grid can vary. Also, other GHGs such as methane may be produced. Consequently, real-time carbon emission data obtainable from the electricity suppliers should be used and this data should take into account other GHG emissions that have been converted into *carbon equivalents* (CO_2eq).

Ideal value of 0.0

Unlike PUE, CUE has dimensions kilograms of CO_2 (KgCO_2eq) per kilowatt-hour (kWh) while PUE is unit-less. CUE has an *ideal value of 0.0*, indicating that no carbon use is associated with the data centre's operations. Both CUE and PUE simply cover the operations of the data centre. They do not cover the full environmental burden of the life-cycle of the data centre and the carbon emissions of the energy consumed in the manufacturing of the IT equipment, the *embodied* energy.

An equivalent but alternative formulation of CUE is the following equation:

$$\begin{aligned} \text{CUE} &= \text{CEF} \times \text{PUE} \\ &= \frac{\text{CO2 Emissions (CO2 eq)}}{\text{Unit of energy (kWh)}} \times \frac{\text{Total Data Centre Energy}}{\text{IT Equipment Energy}} \end{aligned} \tag{4}$$

Carbon Emission Factor (CEF)

Where *CEF* is defined as the *Carbon Emission Factor*, the carbon footprint of
every kWh consumed.

CEF is calculated from data given by electricity suppliers using real-time CO_2 meters (usually on-line and updated every 15 minutes) or government energy agencies.

CUE is a metric commonly used when Sustainability is an issue for consideration in the operation of a data centre. It affords an assessment of the following concerns:
– How a data centre's carbon footprint compares with similar data centres.
– Monitoring the effectiveness of a data centre's Sustainability policy.
– Consideration of the net effect of switching to alternative (renewable) energy sources.
– Comparing the environmental impact of different energy strategies.

Annual CO2 emission tonnage, CO2(Ton)

To calculate the actual *annual CO2 emission tonnage, CO2(Ton)* of a data centre then the following equation is used:

$$\text{CO2(Ton)} = \text{Total Energy Consumed} \times \frac{\text{CO2 Emissions (CO2eq)}}{\text{Unit of energy (kWh)}} \times \frac{1}{\text{(1-DistrLoss)}} \tag{5}$$

DistrLoss

Where *DistrLoss* is the energy lost in transmission from the centres generator or from the grid to the main energy inlet to the centre.

REFLECTION 3

Which energy sources do and which do not contribute to the carbon footprint?
Explain your answer.
What are the pros and cons of these energy sources?

3.9 WATER USAGE EFFECTIVENESS [15]

Water Usage Effectivess (WUE)

The metric *Water Usage Effectivess (WUE)* is defined as:

$$\text{WUE} = \frac{\text{Annual Water Usage}}{\text{IT Equipment Energy}} \tag{6}$$

WUE has dimensions litres/kWh.

WUE_{source}

WUE can refer to the actual water used on-site, primarily that consumed in the cooling operation of the centre, or take into account additionally the water usage involved in the production of the electricity supply $\text{WUE}_{\text{source}}$.

$$\text{WUE}_{\text{source}} = \frac{\text{Annual Site Water Usage} + \text{Annual Energy Production Water Usage}}{\text{IT Equipment Energy}} \tag{7}$$

Energy Water Intensity Factor (EWIF)

To calculate the $\text{WUE}_{\text{source}}$, the water involved in producing each kWh must be known. This can be acquired from the energy supplier. In The U.S various energy agencies use the *Energy Water Intensity Factor (EWIF)* (L/kWh). In this case,

$$\text{WUE}_{\text{source}} = \frac{(\text{EWIF} \times \text{PUE}) + \text{Annual Site Water Usage}}{\text{IT Equipment Energy}} \tag{8}$$

WUE is more easily monitored as it can easily be metered on-site. WUE can be reduced by taking the following actions.
– Reduce IT energy use, so that cooling is reduced and consequently water consumption.

– Optimise the humidity levels of the data centre so that it is running at the low end of the ASHRAE-recommended guidelines for humidity (5.5°C dew point).
– Implement all appropriate best practice airflow management strategies to improve cooling efficiency.
– Operate the data centre at or near the ASHRAE-recommended upper limit for temperature, as this will (depending on the cooling plant) allow warmer chilled water and require less evaporation of water to produce it.

$\text{WUE}_{\text{source}}$ is used when a more holistic environmental impact assessment is required. Reducing water usage in a cooling system may increase its energy consumption, constituting an aggregated increase in water consumption.

The U.S weighted average for thermoelectric and hydroelectric water use is 7.6 Litres of evaporated water per kWh (excluding cooling), this implies that 1 MW data centre has a hydro-footprint of 76 million litres of water/annum or a hydro-footprint of 99 million litres of water/annum including cooling.

4 Monitoring the Data Centre: An Overview

With an active and effective energy management system in operation in a data centre there is the potential for significant energy savings. The savings are achieved by two pervasive trends which emerge as the energy management system becomes more deployed and a mainstream component of daily operations. These trends are:

Technical

1 *Technical* decision making becomes more informed and incisive.
The dynamics and operation of the data centre become more transparent and understood. Information on the interaction of server activity and the power and cooling systems can be reviewed, analysed and predicted. The effects of new strategies, operating procedures or optimisations can be observed and quantified. Overall, technical decisions are more informed and assessment more rigorous and scientific.

2 Working practices, procedures and attitudes concerning energy usage and awareness can be influenced for the better by disseminating information which highlights opportunities for savings and rewards positive *behaviour*.

Behaviour

Annual energy budgets can be established for various organisational cost centres and then monitored on a weekly/monthly basis. The UK's Joint Information Systems Information Committee (JISC), which funds IT in the UK's third-level sector, has reported a number of successful energy management programmes in the UK university sector which has substantially reduced the energy costs and carbon-footprint of several major universities. The University of Manchester is reducing its annual energy bill of £15 million by setting an IT 3% per annum energy reduction target by monitoring daily energy performance of each department. Summaries are displayed in buildings. In the University of Cardiff their high performance cluster has achieved a PUE of 1.3 in part by publicising the energy consumption of different storage configurations. Actions like these are anticipated to make substantial reductions in the UK's third-level education sector bill of £150 million with a carbon footprint of 500,000 tons.

4.1 DCIM, GENERAL REQUIREMENTS

The metrics such as PUE, and CUE that characterise a data centre are single entities. By themselves they give a fairly good picture of the overall energy and carbon credentials of the centre. They require energy consumption data to be measured and aggregated from the IT and non-IT equipment and typically this is performed by the DCIM system of the centre through *sensors* on racks, PDUs etc. Since these are just global, aggregate entities, taking measurements at the main power source of the centre and then measuring the power consumption at the main distribution points to the IT equipment in the centre at the rack or PDU level is sufficient to discern the IT and non-IT and components.

Sensors

As explained previously, taking readings at different times of the day and over an extended period of time to cover different external climatic conditions and loading of the centre, a yearly value for PUE and CUE can be determined. So, in theory the PUE can be calculated fairly straight-forwardly.

REFLECTION 4
What are the major variations in energy consumption in a data centre
– during the day
– during the week
– during the year?

However, knowledge of just these metrics by themselves gives very little insight into the dynamic behaviour of the centre at sufficient resolution, to be able to make informed decisions leading to more optimal use of computing resource at reduced operating costs. Details of server behaviour such as CPU usage and energy/power consumption (amongst other metrics) are necessary to drive and direct efficiencies, and to evaluate the validity and merit of an organisation's energy management programme. There are two main inter-related aspects to energy/power monitoring in a data centre:

Extensiveness

1 The *extensiveness*, frequency and resolution of the energy monitoring system distributed throughout the centre and which feeds the basic energy data into the centre's DCIM system.

Granularity

2 The *granularity* of model of the data centre inherent in processing structure and data-base of the DCIM platform.

Cost
Law of diminishing returns

In the background to all of these considerations is the issue of *cost* and the *law of diminishing returns*. Spending twice as much on monitoring equipment doesn't lead to a centre that is twice efficient. It is essential to ascertain in advance of installing any monitoring equipment, to determine:

What should be measured

a *What should be measured* and why? In a co-location situation measuring power/energy at the rack level may be adequate to check energy bills, for a private cloud with cost centres, measuring power at the server-level may be required.

Level of resolution

b What are the measurement objectives in terms of what is measured and at what *level of resolution* and detail e.g. Power measurement at 90% accuracy at rack, server, or VM level, and temperature and humidity readings at 2% accuracy at server inlets.

Cost

c What is the *cost* of measuring for a given set of sample points and stipulated accuracy. Cost considerations should include the capital cost of the meters, installation costs such as re-wiring and floor retrofitting and downtime and maintenance.

Trade-off between accuracy and cost

d What is the *trade-off between accuracy and cost*. What system or monitoring configuration gives maximum benefits for least cost to the organisation. Flexibility of the monitoring system to be upgraded.

In the data centre, there are three approaches to energy and environmental monitoring:

Physical sensors connected by WiFi

a Traditionally, any form of monitoring involved *physical sensors connected by WiFi* or wires to a central monitoring station. These have a number of problems, their physical imposition, cost of deployment and Wifi systems have transmission blackout regions due to the presence of metal objects in a centre.

Software based

b There has been progression to use more internet *software based* solutions using for instance the SNMP protocol to read electrical or environmental data from devices. These are easier to deploy and in many cases less costly than physical solutions.

Proprietary software agents
Communicate with a DCIM platform

c Installation of non-standard, *proprietary software agents* of IT equipment such as servers which periodically *communicate with a DCIM platform*. The main concern with these approaches is that of security but particularly with banks and financial institute accepting them, this is an indicator that this concern is disappearing.

Simplistic model
Real-time model

The typical architecture of a data centre, the embedded monitoring system and the flow of measurement data to the DCIM platform is displayed in Figure 6. The DCIM system receives periodically (typically of the order of minutes unless there is an anomalous event) updates from the centre. Updates can consist of data regarding energy, power, temperature and humidity or security alerts. Inherent in the DCIM platform is a perceived energy/power model of the centre which must be supported by the monitoring system. It may be the case that a *simplistic model* of the servers is employed in the modelling so that a server is modelled as a device having just 3 or 4 discrete power states. Alternatively, a more accurate *real-time model* may be required of the servers. These are issues explored below, but the decision of which one to use always has a significant cost factor associated with it.

Actions

Regardless of the type of monitoring that is adopted, the data that it generates is transformed by analysis into a number of *actions*. These actions may be taken in real-time affecting immediately the power and cooling operations of the centre. Alternatively, data may be assembled over a extended period of time and used to inform major policy decisions regarding the centre's operation or even for something more practical such as producing bills for users.

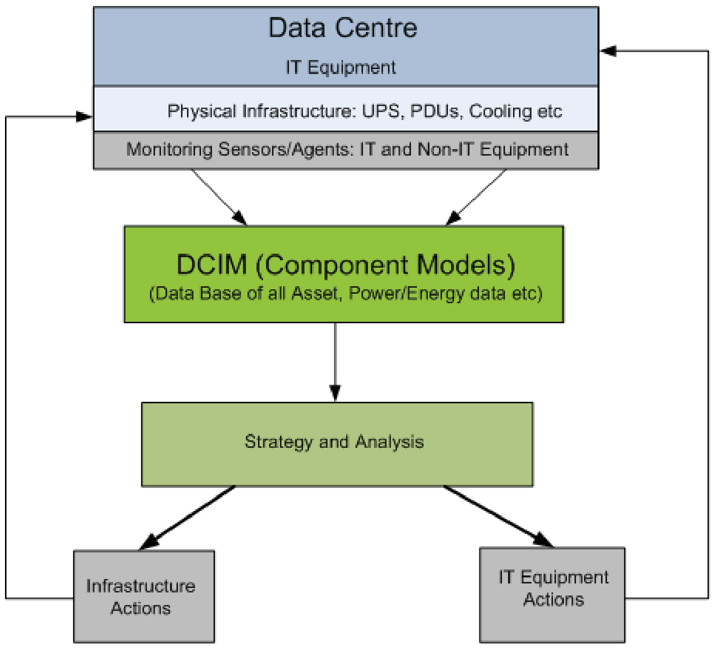

FIGURE 6 The DCIM system is Central in Energy Monitoring and Management

4.2 ENERGY AND POWER MONITORING TECHNIQUES

Infrastructure equipment

Usually monitoring the *infrastructure equipment* is relatively easy, as the number of units involved is a lot less than the number of servers and IT equipment and most of the infrastructure equipment has standard protocol interfaces (e.g. SNMP, Modbus) for reading power/energy etc. which can be accessed easily by the DCIM system. Monitoring the IT equipment is more problematic. There are substantially more IT units in a centre and monitoring them can be more varied. The most difficult to monitor are the servers because there are four broad *technique to monitor server power* each with their own advantages and disadvantages.

Technique to monitor server power

Discrete Server Model

1 *Discrete Server Model*: A server is modelled by running a number of benchmark applications on it and measuring its power consumption with a meter. From this analysis a Minimum (0% utilisation), Average (40-60% utilisation) and Maximum (100% utilisation) rating is produced for a server and sometimes an additional power measurement is taken of the server when it is connected to power but not switched on. This calibration process is done by a number of companies like Power Assure and their PAR^4 certification service and monitoring tools. (PAR^4 refers to their benchmarking process which assesses the power consumption of a server at 4 operating conditions, Plugged-in but not turned on, Idle, 100% CPU utilisation and Maximum power consumption).
The Gartner group also offer a software tool in their RPE2 methodology which gives power ratings of over 25,000 servers and storage devices in different configurations. The power ratings are not derived empirically but rather theoretically from a number of cited benchmarks and relative performance figures.
A common benchmark used in measuring and monitoring server power consumption is the Standard Performance Evaluation Corporation's SPECpower_ssj2008. The benchmark has been designed to exercise a server over an operating envelope which ranges from virtually 0% to 100% CPU usage, giving power values that are quite representative of the server's actual behaviour.

How they will be used

With the models there is still the decision on *how they will be used*. The models require the data on CPU usage and whether this is measured by some user-designed or proprietary code on each server in real-time, or logs on a server are taken each hour, day or week, this is a decision for data centre operations. It should also be recognised that these models are fairly inaccurate for measuring individual server or application power consumption over small time intervals (>15% for 1 minute intervals) and are suitable only for analysing average aggregate server behaviour over relatively long time periods (e.g. 10's minutes).

Physical on-board metering

2 *Physical on-board metering*: Large processor manufacturers such as Intel now have physical meters on their server boards that measure power and temperature. The meters can be accessed through proprietary SDKs such as Intel's RPE2 manager software. Access to the manager is via industry-standard server management protocols such as Intelligent Platform Management Interface (IPMI) or Web-services Management (WS-Web). Intel's DCM manager has been licensed to other commercial players in the DCIM market for integration into their tools. This metering provides accurate real-time data. For various technical reasons this metering is not possible on all server configurations.

Agent-based metering

3 *Agent-based metering*: Whereas physical metering requires meters to be pre-installed during the server specification, agent-based system can be installed at any stage on any server. The agents monitor certain O/S parameters which are periodically communicated to a central master which has a library of accurate power models which use these parameters to calculate the power and energy consumption of the servers. These models offer the greatest degree of resolution offering power monitoring to application and VM(Virtualised machine)-level. Stratergia's Papilon energy monitoring is an example of this approach.

Physical Rack-level metering

4 *Physical Rack-level metering*: Some rack systems are feed by rack-power strips which supply power by cables to individual servers. Each cable supply has a meter which can be interrogated through a standard protocol interface. Rackwise is a company which produces such technology. This approach requires the most physical infrastructure intervention to operate.

Sample servers

In many data centres because there are legacy devices and a mixture of server types, and centres have grown organically over the years there is usually a variety of metering and monitoring methods. In many instances, *sample servers* are measured manually with meters to ascertain the power characteristics of the equipment. In its simplest form, this information can be used in spreadsheets of the data centre assets to calculate different operational scenarios, or it can be used in the DCIM models to give a more automated analysis.

4.3 CARBON-FOOTPRINT ESTIMATION [16], [17], [18]

Calculating the carbon-footprint of an entire data centre, assuming there is no internal power generator using carbon-based fuels, can be calculated from equation (7) above. If, a finer degree of resolution is required, for instance, to determine the carbon-footprint of an individual server then the in-direct energy consumed by ancillary devices utilised by the server have to be considered in addition to the energy consumed by the server it-self. To determine the average energy consumed on the external devices by the server's activity requires a number of bench-marking experiments over time. For instance, the percentage of memory accesses to a memory device by a server can be monitored while the power of the server is observed. In a simple model the server's memory budget is allocated a power consumption value as a proportion of the memory bandwidth. Typical, energy budget overheads for a server are the shown in Table 2. assuming the server power has been scaled to 100. Consumption of ancillary devices has to be taken into account when calculating the individual carbon-footprint of IT equipment such as servers. The % values also indicate the typical power consumption breakdown in an average data centre.

TABLE 2 The Power Energy budget overhead, server is scaled to 100

Server	100 (36%)
Cooling	105 (38%)
Storage/Memory	30 (11%)
Power Losses (Distribution)	25 (9%)
Network	10 (4%)
Lighting	2 (1%)
Ancillary	2 (1%)

Total 274

In this example, if this server was representative of the entire data centre then the PUE is approximately:

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{Total IT Power}} = \frac{274}{(100+30+10)} = 1.95 \tag{9}$$

But the carbon-footprint of the server is not the PUE factor as it induces energy consumption in IT and non-IT equipment. Its energy overhead is 2.74, so that its carbon-footprint for every 1 Whr is actually:

$$= 2.74 \times \frac{\text{CO2 emissions (CO2 eq)}}{\text{Unit of Energy (kWh)}} \times \frac{1}{(1 \text{ - DistrLoss})} \tag{10}$$

Different servers and users have different processing profiles so the overhead analysis has to take into account the specifics of each case.

4.4 COST OF DATA CENTRE DATA COLLECTION

For any energy management system, there has to be some level of power/energy monitoring system. As stated in section 4.1 there is a trade-off between cost of installing a metering framework and the degree of energy resolution required. The most basic objective for any data centre is to determine its PUE and this only requires a broad perspective of how much power is feeding the IT equipment, it is not necessary to know how much goes to each piece of the IT inventory.

REFLECTION 5
Figure 5 gives the heat generation (energy consumption) of the various equipment components in a data centre as a percentage of the overall centre's power consumption. If there is a 10% error in the measurement of each of these entities how do they contribute to the overall error in the estimation of the power consumption of a 1 MW data centre.

Table 3 illustrates the exponential rise in the cost of monitoring a data centre at high levels of accuracy. It demonstrates the approximate cost for a 1 MW data centre (= 1000-2000 servers). It assumes that physical metering on servers is employed where a high degree of power resolution is required and this involves the unit cost of each meter and the labour costs of installation. Costs would be considerably less if software, agent-based metering was applied. The table demonstrates that for virtually no cost an energy model of the centre can be adopted by simply counting the number of servers and multiplying this by an average power value per server. This model can be improved upon by augmenting it with metering values taken from the UOS system which will give some indication of the power distribution to IT and non-IT equipment. This can be improved further by categorising the IT and non-IT equipment and using more detailed energy information to each category. This would be accomplished by processing an inventory data base. The highest level of monitoring giving greatest power resolution on each item of IT equipment is the most expensive and requires extensive monitoring at rack and/or server level.

Level of metering resolution is required

This analysis really invites the question, what *level of metering resolution is required* and how is this cost justified. To answer these points, the objectives and purposes of the metering must be established, and the outcome evaluated in terms of Return on Investment (ROI). This will be discussed in section 4.5.

TABLE 3 The Accuracy of Power Measurement versus Cost

Model	*Level of Power Resolution*	*PUE Error*	*IT Energy Error*	*Metering Cost/1000 servers*
	Server Count	60%	40%	€0
	UPS power monitoring	50%	33%	€0
	Basic Asset monitoring	25%	25%	€0
	Detailed Asset monitoring	15%	12%	Labour cost (€5,000)
Discrete server model	Server classification	12%	12%	Labour + some metering of types (€15,000)
Physical rack-level metering	Rack monitoring	5-10%	5-10%	€100,000
Physical on-board Server (not all types)		2%	2%	€50,000
Agent-based Server/App/Virtual		2%	2%	No Physical cost

4.4.1 *Metering for Basic PUE*

For determination of a data centre's PUE to within approximately 10-15% accuracy, this can be accomplished with minimal metering investment. However, the coarse granularity of the metering will give very little insight and vision of how and where to make changes in operations or technical amendments to improve performance. With a low metering resolution any improvements in the centre would have to be on large scale for any discernible effects to be observed. So if a major cooling or power unit is changed the effect of this would be noticed. In contrast, some significant improvements in centre efficiency can only be achieved at a server-level and then to be observed, it has to be magnified over the entire, or a sizeable proportion of, the server population. This is highly undesirable as any negative effects would possibly have major consequences. It would be preferred to experiment with any new proposals on a limited subsection of the centre's assets and observe any benefits or advantages before exposing the centre to the changes.

Even with knowledge of the potential options for operational or technical improvements, detailed power information is required, which may be elusive with basic power visibility. In general, low metering resolution will limit the options for energy management in any centre.

4.4.2 *High Metering Resolution*

With metering to rack-level and server-level the choices for energy management and energy reduction are wide and various. Of particular importance in energy reduction and only feasible at high metering resolution is *Rack-Power Analysis* and *Server Consolidation*. These are described as follows:

Rack-Power Analysis
Server Consolidation
Rack-power Analysis

Rack-power Analysis: Referring to Figure 3 it was mentioned that stranded power contributed to energy wastage in the data centre and was a consequence of inaccurate estimation of individual server power requirements. By monitoring power at server-level, the power characteristics of each server can be fully known and the total power consumption of the servers in a rack with an adequate safety margin can be factored into the energy requirements. This allows more servers to be consolidated into each rack. This leads to savings in the centre through reduced rack requirements and floor space and more energy efficient PDUs. With an estimate of the capital savings and knowledge of the energy savings that can be estimated with the metering, the overall costs benefits of any proposal can be calculated.

Rack-power Analysis

Server Consolidation: Usually the metering process gives some information on server CPU utilisation, indeed in Stratergia's power model, CPU usage is readily available. Even where this is not possible, generally low CPU usage can be detected as a server with fairly constant idle power consumption (this still can be 50-60% of the server's maximum power consumption). If server has an average (daily/weekly) CPU utilisation rate less than 15% (say), or some other agreed arbitrary threshold, then it is a candidate *for de-commissioning*. Instead of five servers each operating at
5-10% CPU utilisation, it is more energy efficient to have their workload, if possible, transferred to one server which will have a CPU utilisation rate of 50-60% (as cited in Chapter 4 and in section 2 of this Chapter), one server can increase its CPU utilisation rate by 500% with only a 40% increase in energy consumption). This server consolidation doesn't necessitate virtualisation. If servers have the same O/S, providing there are no administration restrictions and licences are transferrable, servers can be *consolidated* quite straightforwardly by shifting one set of applications onto a single server. More consolidation may involve virtualisation. Like rack consolidation, server consolidation will in most cases increase economies in rack, energy and floor space.

For de-commissioning

Consolidated

Detailed power monitoring

In addition to these latter opportunities, only *detailed power monitoring* will facilitate the following energy management options:
– Determine TCO of servers.
– Accurately analyse daily and strategic data centre performance, leading to more efficient capital expenditure and efficient operating costs (e.g.idle times of servers, optimal use of low-cost power tariffs).
– Accurately calculate PUE/DCiE metrics.
– Implement cost centres and charge back for IT services based on actual usage.
– Provide Econometric statistics such as Transactions/Watt or other performance/watt that will help bridge the communication gap between the CFO and CIO perspectives.
– Focused cooling in real-time, to areas of actual power generation.
– Ascertain which servers give optimal power /performance for various applications.
– Use historic data to predict future provisioning.
– Optimise use of data centre real-estate.
– Validate effectiveness of any data centre Green strategy or policy.
– Prioritise in a server replacement programme.
– Simulate of data centre scenarios.
– Calculate of carbon foot-prints per user.
– Analyse feedback of any energy management strategy.

4.5 RETURN ON INVESTMENT (ROI) [19]

In any organisation, while there be concern for the environment and they may have altruistic motives in reducing their energy consumption and carbon footprint, undoubtedly a cost benefit analysis will be required to quantify the net financial gain. This involves re-appraising the Total Cost of Ownership (TCO) of operating the services and IT equipment and seeing the effects of any new efficiencies. The cost of running the IT equipment being a huge component of the centre's operation, estimating the TCO of these entities is central to any assessment. The TCO of IT equipment has directly identifiable line items, its capital cost, annual maintenance and depreciation but it also incorporates the following requirements and cost considerations:
– Total Building Space (Wall space).
– External space (Car park etc.)
– Whitespace (The Raised Floor).
– Critical Load Capacity (The computing payload that the UPS system will support, this is less than the power from the grid.).
– Rack Power density. With higher power requirement from newer server types e.g. blades, not only is the power capacity of the UPS and PDUs a consideration, the location and density of the power demand is also an issue).
– Effective Usable space. Due to power distribution issues, rack consolidation and availability of whitespace not all data centre space is equal. There may be numerically adequate space but not when quality is a factor.
– Direct energy costs.
– Cooling and Power distribution overheads. The capital and running costs of these entities and the net overhead per server. Or IT unit. The Tier rating (1, 2, 3 or 4)of the centre will introduce increasing redundancy and costs into this category.
– Staffing requirements. Any automation of existing processes will have an effect on staffing levels.

3-year period

The TCO is normally estimated for a 1-year period and extrapolated over a *3-year period* (on a linear basis) to give a TCO for 3 years which is the accountancy period at the end of which IT equipment is considered to be obsolete (valueless).

(Accumulated) Net Benefits

Once the TCO of the items involved in the new proposal or action, actual or theoretical is known, the *(Accumulated) Net Benefits* can be determined:

(Accumulated) Net Benefits = Gross Benefits – Ongoing Costs.

Normally this is calculated on an annual basis. Gross benefits are the fixed savings per year as a consequence of the action. It incorporates and financially quantifies all changes in personnel, floor-space, energy reduction etc.

The 3-year return on investment is given by the equation:

$$\mathrm{ROI}_{(3\text{-year})} = \frac{\dfrac{\text{Net Benefit}_{\text{year}}}{(1+\text{Interest rate})} + \dfrac{\text{Net Benefit}_{\text{year2}}}{(1+\text{Interest rate})^2} + \dfrac{\text{Net Benefit}_{\text{year3}}}{(1+\text{Interest rate})^3}}{\text{Initial Investment}} \quad (11)$$

Interest rate

Where *Interest rate* is the rate at which the initial investment would appreciate.

REFLECTION 6

Calculate the Return on Investment for the following scenario:
A new server system costs an initial investment of €90,000, saves €6,000/annum in energy and reduces the IT team by 1 staff member at a saving of €40,000/annum. However, the servers have a maintenance charge of €3,000/annum and cost €6,000 in electricity to operate and all the overheads and depreciations are a further €4,000/annum. What is the Gross Benefit, the Net Benefit the TCO and the Return on Investment over a three year period.

This ROI evaluation allows organisations to predict and evaluate different investment strategies and scenarios before committing to action.

5 Case Study: Analysis of an Actual Energy Monitoring and Management System- Papillon

5.1 PAPILLON INTRODUCTION

Non-intrusive

Agents

This is case study of an energy monitoring and management system for data centres designed by Stratergia Ltd (Ireland). The PAPILLON system uses a client-server type architecture which is totally *non-intrusive* and benign to the performance of the data centre's servers and communication network. *Agents* installed on each server monitor and acquire server behaviour data in the background. Periodically, data is communicated to a master server which has power models for each server type in the network. Using the behaviour data with the power models, the master computes and saves to a database the power consumption and other power related information for each server. This forms the foundation for the large portfolio of PAPILLON power analysis tools. Some of these tools highlight which energy-saving actions can be taken and quantify the energy that will be saved.

Power models are generated using benchmarks which comprehensively exercise the server. While the benchmarks execute server power is measured periodically with a meter. By combining the meter values with the aggregated operating system parameters the power behaviour of the server can be mathematically modelled. Every server type in a data centre has a unique power model generated in this process. This process is shown in Figure 7 and can be performed on-site or by Stratergia.

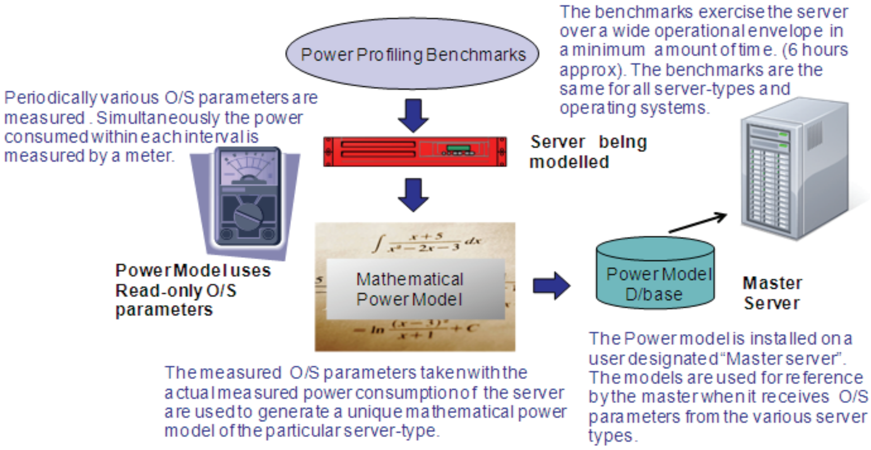

FIGURE 7 The Server Power Model Generation Process

The installation of the monitoring agents is performed over the internet, in contrast to physical meters, there is no down-time or a requirement for any rewiring or retro-fitting of data centre resources, see Figure 8.

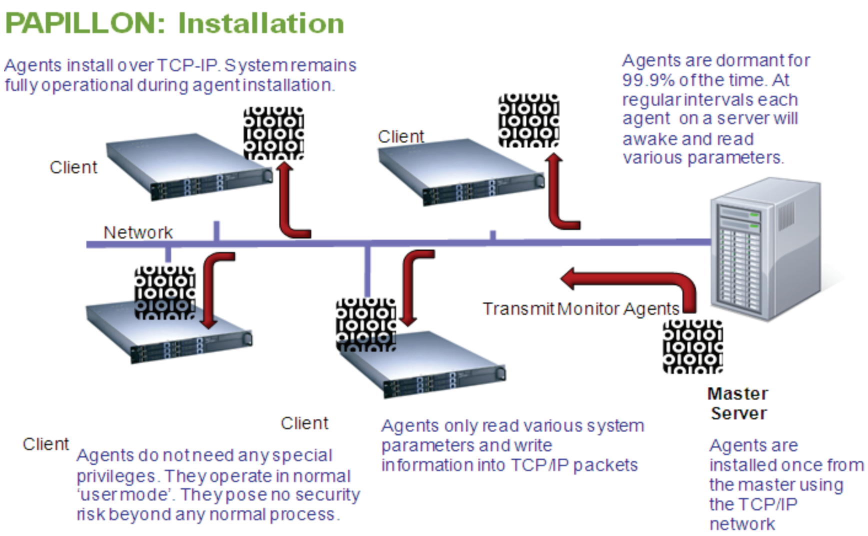

FIGURE 8 Installation of Agent Monitors Circumvents the Need for Physical Meters

Master server

Once installed, clients periodically transmit specific server activity data, acquired through the O/S, to a designated *Master server*. The Master uses its library of power models to compute the power and energy consumption of every server in the centre, see Figure 9. It also receives information from each server concerning the processes that are currently executing on them. With this information a comprehensive, real-time overview of the power demand of every process, server, rack and the entire data centre can be maintained in the Master's data base. The data base can be accessed through an extensive API library by a suite of open source or proprietary analytical tools. Amongst the hierarchy of metrics and statistics generated by the tools and presented on dashboards and reports, are indicators proposing and quantifying energy-saving actions such as rack and server consolidation.

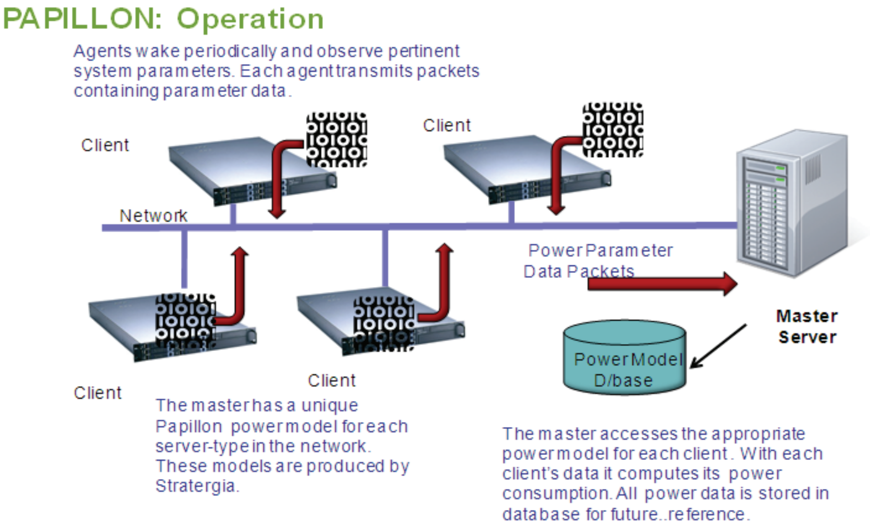

FIGURE 9 Periodic Metering of Servers through Agent/Master Dialogue

Actual dashboards from the Papillon system, their purpose and interpretation, are presented below in Figs 10 to 13.

5.2 PAPILLON ENERGY DASHBOARDS

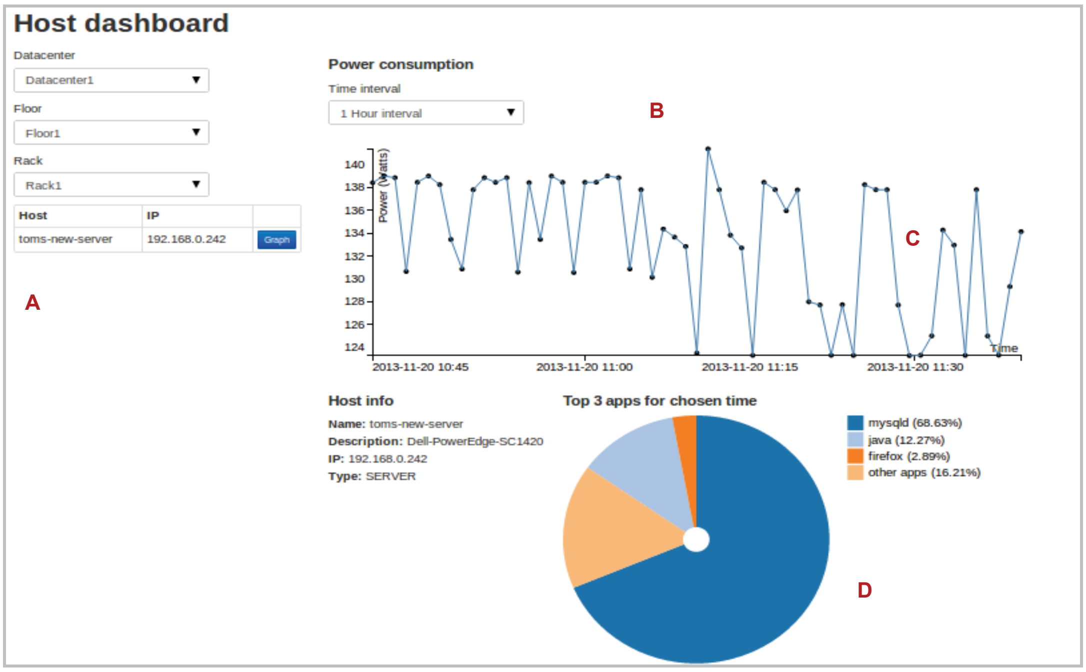

Legend

A: Descending the designated data centre, floor, rack hierarchy via a number of pull-down menus a specific server is selected.

B:The sample period to be displayed in the power graph is specified.

C: The power behaviour of the selected server for a given time period is displayed. By clicking a point on the graph, a breakdown of the power consumed by the active processes in this interval is displayed in D.

D: The power distribution shared by the various processes in the chosen time interval is displayed.

FIGURE 10 The Real-Time Application Power Consumption Dashboard in Papillon

Inspected in real-time

Purpose for Dashboard: Any server can be *inspected in real-time*, so that if there is a suspicion of anomalous activity, it can be quickly investigated. Furthermore, the power consumption of individual pplications or users can be monitored with the intention of identifying which servers give best performance per watt for a given loading in a particular application.

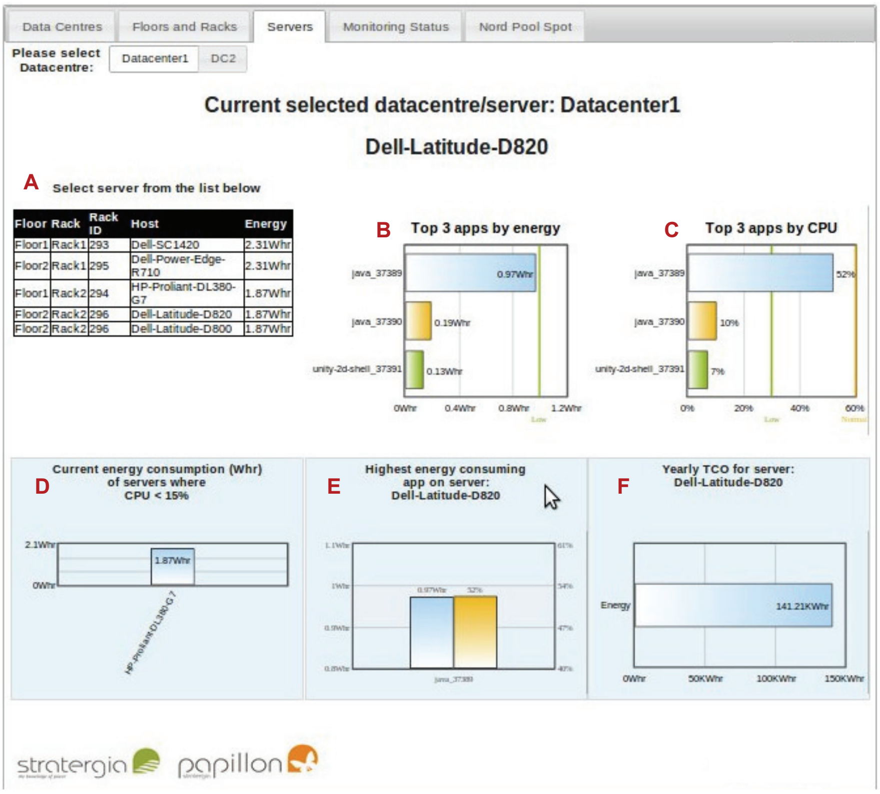

Legend

A: A particular server is selected via drop-down menu.

B: The top three applications with the largest energy consumption over a given period (normally a week) are identified together with their consumption data.

C: The top three most CPU intensive applications on the server are ide

D: For the rack on which the current server is resident, other servers which are under-utilised (< 15% usage) are highlighted together with their power consumption.

E: The top consumer application statistics is displayed, energy consumed and energy consumed as percentage of overall consumption.

F: The yearly TCO of the selected server based on its measured energy consumption over the given period.

FIGURE 11 The Server Power Consumption Dashboard in Papillon

Effectively used

Purpose for Dashboard: The dashboard presents a range of energy related data which is very useful in seeing how a server is being *effectively used*. The server's CPU usage together with a breakdown of the largest application consumers indicates how it is being used and by which applications. The server can be identified if it is a candidate for decommissioning by its inclusion or absence in the list of low activity servers for the associated rack The two statistics shown in the dashboard, TCO and CPU taken collectively, is a good indicator of the productivity of the server.

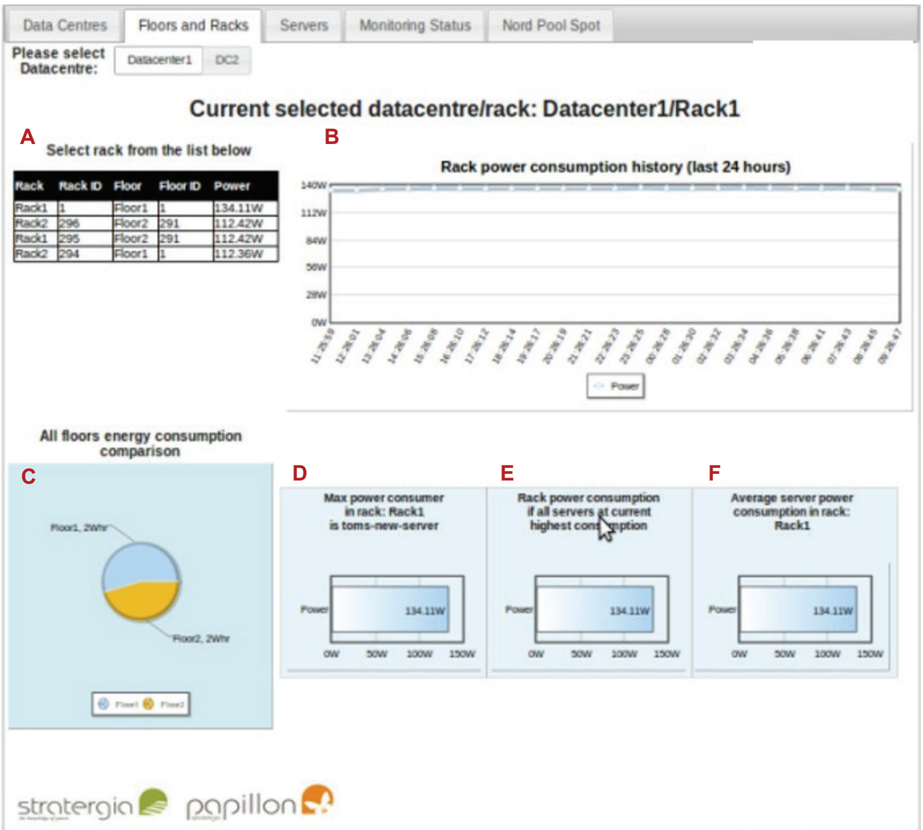

Legend

A: Racks are displayed together with their location and power consumption averaged over a given period. From this list a particular rack can be selected for analysis on the dashboard.

B: The power consumption over the last 24 hours of the selected rack is displayed.

C: The energy consumed on each floor of the data centre over the last 24 hours is displayed.

D: The power of the biggest server power consumer for the given period in the rack is identified.

E: The worst case power consumption of the rack is calculated and displayed. This scenario depicts the situation where all servers simultaneously consume their maximum power consumption as detected for each server over the given period.

F: The average power consumption of the server over the given period.

FIGURE 12 The Rack Power Consumption Dashboard in Papillon

Power performance

Purpose for Dashboard: The dashboard allows enquiry into the *power performance* of a rack and its candidacy as a rack for increased server residency. Alternatively, racks will be detected that are close to their maximum power ceiling and which maybe should be de-populated of servers.

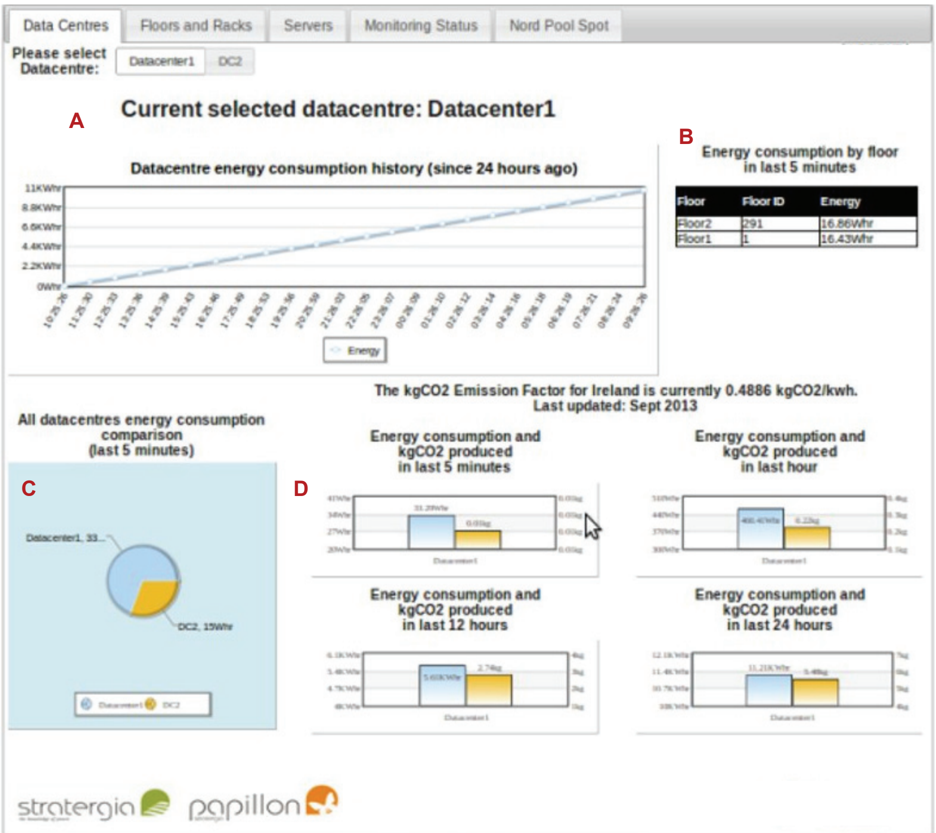

Legend

A: There may be several data centres (or sub-sections) in a organisation. This dashboard permits a particular data centre or sub-section to selected and its energy consumption over the last 24 hours is displayed.

B: The energy consumption of the selected centre over the last 5 minutes is displayed.

C: The energy consumption relative to other data centres or (sub-sections) in the organisation is displayed.

D: The energy consumption and carbon footprint of the centre over the last 5 minutes, hour, 12 hours and 24 hours is displayed.

FIGURE 13 The Overall Power Consumption Dashboard in Papillon

Overall

Purpose for Dashboard: The dashboard gives visibility of the *overall* energy and carbon footprint status of the data centre. It is particularly effective in observing the immediate centre-wide effect of any new operations, technology or procedures that may have been introduced.

5.3 PAPILLON REPORTS AND IDENTIFICATION OF ENERGY SAVING ACTIONS

Energy-saving actions

While dashboards are excellent channels for exploring the hierarchy of the data centre and permit investigations to drill-down and resolve problems, or reviews which lead to a better understanding of the anatomy and behaviour of the centre and which can be incorporated into strategic updates, they are not a practical means for identifying *energy-saving actions* or threats to the operational viability of the centre.

Reports or alerts

Due to the vast amount of IT and non-IT equipment in a centre, the amount of information that has to be processed to formulate an understanding and insight of where energy-saving actions and initiatives could be applied, is virtually impossible for a single human. It really requires the energy system of the DCIM platform to automatically process the energy data from the centre and notify or suggest energy-saving actions to the data centre manager in a series of *reports or alerts*. Some of the Papillon reports that may be requested are shown below.

5.3.1 *Papillon Reports*

REPORT 1 Operational and Apparent Wastage for all Servers

Server-ID & IP-Addr	*av_cpu_usage (av_cpu_wk) (val □ 100)% (val ~ 0 -> 1)*	*Energy Consumption for 3 years (eng_cons_3) (kWh)*	*Operational Cost for 3 years (op_cost_3) (€)*	*Apparent Cost Wastage (app_waste_cost_3) (€)*	*Carbon Footprint KgCO$_2$*
N1-admin 192.168.12.234	3%	9,636	1,734	1,682	4,541
N2-sales 192.168.12.237	3.5%	8,900	1,602	1,545 (1602 x 0.965)	3,781

Low-usage are list at the top

Wastage metric Apparent Cost Wastage

Report 1 lists all servers in a designated area that have been automatically surveyed over a specified period regarding their average CPU usage and operational cost. It lists the servers in ascending order of usage, so that *low-usage are list at the top*. This analysis takes the TCO of each server and multiplies it by the percentage of time that it is idle, the result is a *wastage metric*, *Apparent Cost Wastage*. This highlights servers which are candidates for de-commissioning and the financial and carbon savings that would be saved if this action was taken.

REPORT 2 Least & Most Used Server Analysis

Server-ID & IP-Addr	*av_cpu_usage (av_cpu_wk) (%)*	*Top 3 Apps*	*Apparent Cost Wastage (app_waste_cost_3) (€)*	
N1-admin 192.168.12.234	3%	Youtube Sage facebook	1,682	Top 10 Least used servers
N2-sales 192.168.12.237	3.5%		1,545	
N3-RnD 192.156.12.256	85%	Oracle Matlab python	23	Top 10 Most used Servers
N4-marketing 192.168.12.367	63%		36	

Reports 2 is similar to Report 1 but it focuses on very high and very low-usage servers. These are servers which either are close to maximum capacity and which should shed some of their workload, or alternatively very under-utilised and should be de-commissioned.

REPORT 3 Candidates for Server Consolidation

Server-ID & IP-Addr	*av_cpu_usage (av_cpu_wk) (45% to 60%)*			
N1-admin 192.168.12.234	46%			
N2-sales 192.168.12.237	51%			

Report 3 highlights servers which are between 45 to 55% utilised. These are servers that should be considered to accept applications from other servers that are being de-commissioned.

REPORT 4 Candidates for Rack Consolidation

Rack 22

PDU Max (10KW)	
8.1 KW	Worst Case Power
7.2 KW	Actual Max Power
5.3 KW	Average Power

Report 4 gives power statistics on all selected racks indicating the amount of head-room that exists based on the racks actual power consumption. Taken in conjunction with the racks PDU rating, it indicates racks that are fully packed and those that have spare capacity

6 Example Cost Savings Analysis for an Energy Management System

Reflection 7 is a cost-saving analysis for a hypothetical energy management system which outlines the financial benefits that would be accrued if the system was adopted and installed. It only refers to direct savings. Other savings would be specific to the data centre where it is to be installed. These cost savings would include possible man-power reductions/productivity, reduction of software licences changes in work practices and tax credits.

REFLECTION 7
Write a report to justify the expenditure of a new energy management system using the following data and assumptions:
a The average server power consumption is 300 W per server.
b The cost of 1kWhr is €0.20.
c The data centre has a PUE = 1.7.
d Up to 25% of the servers can be detected by the energy management system as functionally redundant and can be decommissioned.
e The footprint of a rack in a Data Centre costs approx. €10,000 to build and provision.
f The energy management system prevents one data centre crash/annum.
Make the case for the system under the following headings by quantifying the benefits in monetary terms:
- I Increased Reliability and uptime.
- II Server Capital Savings and Energy Reduction.
- III Real Estate Cost Savings.

7 Energy Efficiency Method and Control [19], [20]

Having given you the total overview of the configuration of the data centre (Chapter 2), the IT equipment (Chapters 3 and 4) and the points to focus during measurement, we now give a summary of the energy reduction techniques that can be applied in the data centre. Here are some of the major energy efficiency techniques.

7.1 IMPROVE AND MANAGE AIRFLOW

Removing heat from a data centre is equivalent to removing exhausts from a combustion engine. Heat is a by-product of the centre's operation and it must be removed and expelled to allow the centre to function. It costs money (the energy consumed by the CRAC and chiller systems) to remove heat, and therefore the airflow to remove energy should be focused where the heat is being produced in the IT equipment, particularly hotspots. Conversely, areas that are being chilled too much should be avoided. In either case the data centre must have containment areas to section off air flow to designated areas of the centre. Detection of these hotspots or areas of high heat production can be detected by simulation using Computational Fluid Dynamics, thermal imaging, temperature

sensors or monitoring the energy being produced by the servers and other IT equipment. Simple measures can improve airflow and efficiencies such as:
– Checking that the airflow in the server aisles and raised floors are not restricted by equipment or cables.
– Ensuring that panels on the racks and aisles are not missing, so as to prevent hot air mixing with incoming cool air. In general, never let hot and cold air mix. This principle manifests itself in the technique known as Hot Aisle/Cold Aisle containment and reduces operational cooling cost by up to 20%.

Hot aisle/cold

A *hot aisle/cold* aisle containment configuration arranges racks of servers so that the cold air inlet sides of two rows face each other with the hot discharge sides facing towards the hot discharge of the next row. This creates cool air supply areas for intake with alternate row areas that become hot to optimise hot air circulation collection and return for re-cooling, thus avoiding hot air and cold air intermixing. Cold air must be delivered to cold aisles and hot air extracted from hot aisles. In air –cooled racks the chilled air moving over the servers is only a fraction of the total airflow. The remaining air comes from the ambient air supply of the external room or air that is re-circulated. In this air management technique, the dilution effect of recirculation can cause the air temperature at the server inlets at the base of the rack to range from 10oC to 15oC, while the inlet temperature at the servers at the top of the rack may range between 30oC to 40oC. Since the rack heat load is determined by the hottest region, this temperature differential can have a severe limitation on the rack server capacity. There are several methods to resolve this problem, implementing physical barriers to reduce hot and cold air intermixing is one approach, while increasing the temperature differential between hot and cold aisles (this is called *increasing Delta T*) thereby increasing the efficiency of the heat extraction process is another common approach. More advanced techniques, use water-cooled heat exchangers on the racks. Since water has between 50 and 1,000 times the capacity to remove heat than air, such systems have been reported to remove 60% of heat from high density racks (33kW).

Increasing Delta T

7.2 RAISE OPERATING TEMPERATURES

ASHRAE's 'Thermal Guidelines for Data Processing Environments' (2012) recommends a temperature range of 18–27 °C (64–81 °F), a dew point range of 5–15 °C (41–59 °F), and a maximum relative humidity of 60% for data centre environments. In view of these recommendations, since lower operational temperatures require more expensive cooling, the upper temperature range 27°C is now typical used. If possible, provided that the IT equipment can tolerate it, the temperature can be raised to cut power consumption and raise efficiency and some practitioners are pushing the operational envelope to 32°C. These guidelines depend on the elevation of the data centre. Higher elevations require lowering the maximum dry bulb temperature 1°C for every 218m above 1,287m.

7.3 ECONOMISERS

Economisers are mechanical devices used in data centres to support or replace the CRAC and chiller systems by using the cooler ambient air. This potentially can reduce the centre's energy consumption by up to 60%. Since it is dependent on an external air stream that is quite cooler than the centre's computer room temperature, economisers are only useful in cool climates such as Ireland, UK and Scandinavia. Economizers recycle energy produced within a system or leverage environmental temperature differences to achieve efficiency improvements. The outside air must be filtered to remove any pollutants or particulates and its relative humidity must be restricted to between 40% and 55%.
There are two versions of the device used in data centres: air-side economizers and water-side economizers.
– Airside economizers pull cooler outside air directly into a facility which is subsequently heated by the equipment and expelled.
– Water-side economizers use cold air to cool an exterior water tower. The chilled water from the tower is then used in the air conditioners inside the data centre instead of mechanically-chilled water, reducing energy costs. Water-side economizers often operate during night-time to take advantage of cooler ambient temperatures.
In climates where economisers are deployable, they are an integral part in sustainable, green computing best practices.

7.4 POWER DISTRIBUTION

UPSs and PDUs have a significant effect on data centre efficiency. Any loses in these units are generated as heat, so there is a double effect in any loses, the cost of the loss in terms of paying for energy that is not utilised and then having to pay extra for cooling this energy loss. These losses are incurred whenever there is any AC to DC conversion or vice versa or whenever the voltage is increased or decreased. There are many types of UPS systems, double conversion, delta conversion and rotary/flywheel designs. While flywheels are most efficient, most UPS systems operate from batteries. The Efficiency/Load graph for both UPSs and PDUs is very similar to that shown in Figure 4. UPD and PDUs that are 95%+ efficient are now available at 30% loading, but the return on investment compared to ordinary units less efficient servers may be 5 years or more. Less expensive units only achieve this efficiency at loading in excess of 60%.

8 Best Practices for Reducing Energy Consumption and Producing More Environmentally Sustainable Data Centres

These guidelines are based on the directives of the European Code of Conduct on Data Centres and various reports [21],[22].

1 Implement a comprehensive corporate/organisational sustainable energy plan which monitors energy consumption and incentivises users to reduce their consumption. An example is a charge-back mechanism whereby managers of business units are given credit for un-used carbon or energy budgets. Use zero or low carbon-energy where possible in the energy provisioning of the centre. This aspect is quite specific to the locality and geography of the data centre site.

2 Gather sufficient data to generate accurate and current KPIs like PUE, CPU usage, TCO etc. that can suggest and sign-post energy-saving actions and cost reductions.

3 Introduce a virtualisation policy to increase the utilisation rate of servers. The utilisation rates and energy consumption should be closely monitored in order to orchestrate maximum return from this investment.

4 Right-size power and cooling capacity to the demand of the IT equipment. This involves maintaining an efficient and accurate asset tracking system which can indicate the location and details of all IT assets, and an effective and sufficiently accurate power monitoring system encompassing IT, power and infrastructure equipment. This is really addressed by the data centre's DCIM system.

5 In selecting IT equipment chose items which have been endorsed by various standard bodies as being energy efficient e.g. the EU Energy Star (http://www.eu-energystar.org/en/index.html) rated equipment.

6 Invest in training staff to be knowledgeable and up-to-date in the latest server, cooling and data centre technologies and best practices so that these can be competently introduced when required.

7 Analyse the behaviour of servers and their applications/users in order to identify inefficiencies which may be removed by consolidating the number of application licences and/or servers. Analysis may also highlight the option to re-schedule applications which may balance the loading in the centre or enable applications to run at off-peak periods which are cheaper and frequently have a lower carbon-footprint per kWh. Information may also be acquired which indicates that certain server architectures are more energy efficient for running specific applications compared to others.

8 Standardise server and infrastructure equipment as far as possible. This leads to less variability in operation and easier understanding of the dynamics of the centre. Planning and provisioning is more manageable.

9 Evaluate servers for cost/performance (e.g. Transactions/watt) with real-life applications or pertinent benchmarks to ascertain their true TCO and to guide decisions in any procurement replacement strategy.10. Focus on reducing costs and energy consumption on the big energy-consumers, the IT equipment, Cooling equipment and UPS units in the centre are the major players in this area. The ASHRAE 9.9 Guidelines (2008) on data centre temperature and humidity management broaden the acceptable temperature range for data centres to 64.4 to 80.6 degrees Fahrenheit and recommend that the point of measurement for temperature is the air inlet of the IT equipment instead of the room temperature, for high-density racks more measurement throughout the rack should be taken. The humidity range is also extended in this guideline. The purpose of expanding the thermal envelope is to operate the data centre in climatic conditions which require less cooling and consequently less energy. A ramification of operating a hotter environment is that air-flow must be more efficient to avoid hot-spots, and if the UPS is in the same temperature space, the batteries can have a shorter life-span.

SUMMARY

Data centres make a heavy demand on many electrical grid systems, with a significant impact on the environment. From an operational perspective, the energy component is the biggest cost factor and loss or impairment of continuity of supply the most potent risk factor. This chapter explained the driving forces ramping up and accelerating this demand, and how despite the advances in server technology , energy delivery, and cooling, by themselves or even collectively, they are inadequate to address the operational challenges, unless energy measurement and management are given equal priority to any other commercial or financial consideration in the data centre. It was seen that comprehensive energy management is cost-effective, with net gains for the centre's bottom-line (net profits), a more sustainable business model through lower energy requirements, and less impact on the environment as a consequence of smaller carbon emissions and water cooling. Standard metrics such as PUE were reviewed and shown to be deficient when energy efficiencies are part of the analysis. In fact, it will be apparent that due to the complex nature of how energy is distributed and consumed in a centre, no single metric or parameter is adequate in expressing the overall data centre energy performance or efficiency such that it can be used to guide or direct attention where energy-savings can be made. To achieve this, a group of metrics and measurements at server and rack-level must be adopted and which can only be generated by an automated real-time energy monitoring system.

Literature

[1] Mayer-Schonenberger, Cukier: Big Data, A Revolution That Will Transform How We Live, Work and Think. John Murrau Publishers, 2013.
[2] Dalton, Bristow: e-Infranet: Green Sustainability Policies for e-Infrastructures: EU Policy document, Brussels, May, 2012.
[3] Ponemon Institute National Rept on Data Center Downtime. Emerson Network Power, October, 2010.
[4] Stansberry, Kudritski: Uptime 2012 Data Center Industry Survey.
[5] Kaplan, Forrest, Kindler: 'Revolutionizing Data Center Energy Efficiency', McKinsey Report, July 2008.
[6] Reclaiming the Hidden Capacity in Your Data Center: Power Assure White Paper, 2010.
[7] Electrical Efficiency Measurement for Data Centers. APC/Schneider White Paper No. 154, 2008
[8] Belady, Rawson, Pfleuger, Cader: Green Grid Data Center Power Efficiency Metrics: PUE and DCIE. White Paper No. 6. The Green Grid, 2008.
[9] Azevedo, Cooley, Patterson, Blackburn: Data Center Efficiency Metrics: mPUE, Partial PUE, ERE, DCcE. www,thegreengrid.org, 2011.
[10] Longbottom: Musings on Data Centres. Quocirca, ComputerWeekly.com, January 2013.
[11] Combining PUE with Other Energy Efficiency Metrics: Power Assure White Paper, 2011.
[12] Neudorfer: Total Cost of Ownership. Data Center Knowledge, June, 2012.

[13] Belady: Carbon Usage Effectiveness (CUE): A green Grid Data Center Sustainability Metric. White Paper No.32, The Green Grid, 2010.

[14] Azevado, Belady, Patterson, Pouchet: Using CUE and WUE to Improve Operations in Your Data Center. www.thegreengrid.org, 2011.

[15] Patterson, Azevado, Belady, Pouchet: Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric. White Paper No. 35, The Green Grid, 2011.

[16] Longbottom, Tarzey: Managing carbon reduction across your data centre assets. Quocirca, November, 2009.

[17] Rasmussen: Allocating Data Center Energy Costs and Carbon to IT Users. White Paper No. 161, Schneider Electric, 2011.

[18] Bouley: Estimating a Data Center's Electrical Carbon Footprint. White Paper No. 66, Schneider Electric, 2011.

[19] Minas, Ellison: Energy Efficiency for Information Technology. Intel Press, 2009.

[20] Neudorfer: Energy Efficient Cooling. www.searchdatacenter.com, 2011.

[21] Longbottom: Powering the data centre. Quocirca, SearchVirtualDataCentre.co.uk, 2012

[22] Best Practices in Data Center Power Management. Enterprise Management Associates, June 2010.

MODEL ANSWERS

Answers to Reflection Questions

1 The motives are: use as less energy for data processing (lower costs and better availability of power from the grid)
The issues are:
prevent oversizing of the data centre
prevent heat generation
use low energy equipment (Chapters 2, 3 and 4)
finetune the operating system of the IT equipment (Chapter 4).

2 In the data centre example shown in Figure 5 the values of the PUE and DCiE are:

PUE = 100/40 = 2.5
DciE = 40/100 × 100% = 40%

From this example, it is apparent why DCiE is often referred to as the *Efficiency* of a data centre.

3 Energy sources, CO_2 production, pros and cons:

energy source	*CO_2 production*	*reliable energy production*	*remarks*
sun	no	no	amount depends on weather and time of the day
hydro	no	yes	sometimes harm to landscape
wind	no	no	amount depends on weather and time of the day
coal	yes	yes	
gas	yes	yes	
nuclear	no	yes	severe problems in operations (Fukushima) and storage of nuclear material

4 The major variations in energy consumption in a data centre:
– during the day: The work load is mostly higher during office hours
– during the week: The work load is mostly higher during office days
– during the year: The cooling capacity is higher in the summer than in the winter.

5 Total Data Centre Power Consumption = Power Consumption of (IT + Chiller + UPS + + CRAC + PDU + Humidifier + Lighting/Switch-gear)

The error in the power estimate of each component is as follows:
IT: 400 +/- 40 kW
Chiller: 280 +/- 28 kW
UPS: 160 +/- 16 kW
CRAC: 70 +/- 7 kW
PDU: 50 +/- 5 kW
Humidifier: 20 +/- 2kW
Lighting/ Switch gear: 20 +/- 2kW

Measuring the error in any of these entities is independent of the error in any other component (i.e. they are orthogonal). Statistically then, the error in the overall power estimation of the centre using the power measurement of each of these entities is the square-root of the sum of the squares of the error.

Centre power error = [402 + 282 + 162 + 72 + 52 + 22 + 22]1/2
= 52.2 kW

This calculation illustrates the futility of being very precise in the measurement of factors which contribute insignificantly to the overall result. For instance, a 100% error in the humidifier estimate would only increase the overall error to 55.8

Therefore when considering the cost of metering, the effect of accuracy in the meters on the various pieces of equipment on the overall result should be taken into account on a return for value basis.

6 The 3-year return on investment is given by the equation:

$$\mathrm{ROI}_{(3\text{-year})} = \frac{\dfrac{\text{Net Benefit}_{\text{year}}}{(1+\text{Interest rate})} + \dfrac{\text{Net Benefit}_{\text{year2}}}{(1+\text{Interest rate})^2} + \dfrac{\text{Net Benefit}_{\text{year3}}}{(1+\text{Interest rate})^3}}{\text{Initial Investment}} \tag{12}$$

Where *Interest rate* is the rate at which the initial investment would appreciate.

The Gross Benefit is €46,000

The Net Benefit = €46,000 – €13,000 = €33,000

Using the latter example and assuming the savings and costs are the same for year$_2$ and year$_3$ and interest rate is 10%.

$$\mathrm{ROI}_{(3\text{-year})} = \frac{\dfrac{€33,000}{(1+\text{Interest rate})} + \dfrac{€33,000}{(1+\text{Interest rate})^2} + \dfrac{€33,000}{(1+\text{Interest rate})^3}}{90,000} = 91\% \tag{13}$$

This indicates that a return of 91% has been realised on the initial investment allowing for the appreciation that the investment would have gained through normal interest appreciation.

7 The calculation for the report is as follows.

I Increase reliability/uptime
In section 1.3 we have seen that there is on average 1 period of downtime per year due to power and cooling issues. Average recovery time is 134 minutes and costs \$505,000 = (€390,000 approx).”

A data centre of 10,000 servers that saves just 1 downtime/year will save €39 per server/ annum.

Assuming prevention of just 1 downtime saving/annum on 10,000 servers, this is equivalent to a saving of €39 per annum saving/server.

II Server Capital Savings and Energy Reduction
A typical server shows 300 watt power consumption:

In 1 year this server consumes 365 × 24 × 300 Whr
= 2628 kWhr
= €525 assuming the energy price is €0.20 per kWh

With a PUE = 1.7 this has an associated overhead, 0.7 × €525= €368
Total operational cost: €893
In a server population of 1000 servers, 250 are found to be redundant.

The consolidation process is assumed to be simply one of taking applications off servers and installing them on the other servers to increase their utilisation. No virtualisation cost is involved.

The savings of this action are:

Capital Cost saving in servers
250servers × €1000 (Av. cost of a server) = €250,000 per annum/1000servers.
= €250/server.

This will be amortised over a 3 year period, the standard depreciation period for IT equipment

= €83/server/annum

Energy Savings due to Server Decommissioning
250 servers × €893 (Energy + Cooling cost) = €223,250 Energy reduction.

Assuming the computational load of the 250 decommissioned servers can be distributed among the remaining 750 servers with a negligible increase in their power consumption.

The Saving = €223,250/1000 servers = €223.23/server.

III *Real Estate Cost Savings*
The footprint of a rack in a data centre costs approx. €10,000 to build and provision.
A reduction of 250 servers saves approximately 25 rack areas equivalent to a cost reduction of approx. €250,000 per 1000 servers in real-estate. This equates to €250/server. This is once-off saving but can be amortised over 3 years (€83) the normal period before data centres have to be extended.

Conclusion
Total, direct immediate savings per server are:

Reliability/Uptime	€39
Capital Cost	€83
Energy Reduction	€223
Real Estate Saving	€83
Total	€428/server/annum

The energy management system saves €428/server/annum and can be justified economically providing its cost does not exceed this amount per server. Obviously, in practical terms the cost of the energy management system would need to cost substantially less than this amount to be financially attractive.