Saving Energy in the Data Centre Industry

White Paper

Michael Rudgyard
Concurrent Thinking Ltd
March 2014

Potential for energy savings in a typical data centre

There are a number of ways that a typical data centre is able to achieve energy and CO2 savings using Data Centre Infrastructure Management (DCIM) software. Typically, these savings are made through a process of iterative improvements following the identification of potential savings in the areas detailed below.

Base-lining power usage in the data centre

The first step to any data centre efficiency programme is to understand the total energy used by the data centre, as well as how this fluctuates over time, and therefore in line with demand. Without this it is not possible to base-line potential savings, nor calculate how any potential changes might be made given the changing demand.

The second step is to understand the total power used by IT equipment as well as the overheads of cooling and power losses. This can be used to define the Power Usage Effectiveness (PUE) of the data centre, a performance metric that was proposed by the Green Grid organisation in 2006 as:

$$\text{PUE} = \frac{\text{total power used by the data centre facility}}{\text{total power used by computing systems within the facility}}$$

In an ideal world, there would be no cooling or power distribution overheads and PUE=1. However, Kooney estimates that the average PUE is actually between 1.83 and 1.92. He notes that while some of the hyper-scale data centres such as those owned by Google, Amazon and Microsoft can achieve PUEs between 1.1 and 1.2, these companies are in the minority.

A recent survey conducted by Digital Realty Trust indicated that only 20 percent of the 300 North American data centre companies with revenues of at least $1 billion and/or more than 5,000 employees have a PUE below 2.0. The average PUE was discovered to be 2.9.

PUE is now widely adopted within the data centre industry as a means to measure data centre efficiency, although as detailed below, it is largely a measure of cooling efficiency (even if it includes other overheads relating to such as power losses and lighting).

We note that subscribers to the European Code of Conduct for Data Centres, a voluntary code for best practise within the European data centre industry, are obliged to monitor and publish their PUE figures in order to self-certify themselves as participants to the scheme.

Monitoring and managing the efficiency of cooling equipment

The efficiency of air conditioning systems generally reduce significantly over time, despite the extensive and expensive support contracts that are generally in place. However it is rare for data centres to monitor the efficiency of these systems in a proactive manner, and the support contracts are generally only called upon if there is a failure or major problem.

By monitoring the power used by cooling systems, as well as the inlet and outlet temperatures and the air flow rate, adverse changes in the efficiency of systems can be identified and resolved at little or no cost based on existing maintenance contracts.

In many cases, existing problems can also be recognised as soon as a monitoring capability is put in place. For example, a common problem is due to the incorrect commissioning of air conditioning, when cyclic behaviour and over-cooling is evident because of the use of localised control systems foreach cooling unit.

Environmental management

It is often the case that data centres only monitor ambient temperature through building management systems at a small number of sample points. Combined with the conservative nature of the industry and the need to guarantee up-time for mission-critical applications, means that most data centres are massively overcooled.

A 2013 survey by the Uptime Institute found that almost half of the 1000 data centres interviewed were operating at between 71°F (21.6°C) and 75°F (23.9°C), while 37% were operating at 65°F (18.3°C) to 70°F (21.1°C).

This compares with advice from the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) who set the recommended high end temperature from 77°F (25°C) to 80.6°F (27°C) for Class 1 data centres, while the accepted high end was defined to be 89.6°F (32°C) . The survey was voluntary and therefore self-selecting; it is probable that this sample represents data centres with over average efficiency.

To put this into perspective, the U.S. Department of Energy's website estimates a 1% energy saving for each 1°F that the air-conditioning temperature is raised. The U.S. General Services Administration has said that it can save 4% to 5% in energy costs (for the cooling component) for every 1°C increase in the server inlet temperature.

The best practise is to monitor detailed environmental information using multiple temperature sensors within each rack, as well as sensors built into IT equipment. This allows the customer to identify and deal with hotspots: by using blanking plates to reduce recirculation zones, by improving airflow within floor voids, or by redistributing IT equipment within and across racks. Once hotspots have been removed, the average temperature of the data centre can be increased.

Detailed monitoring at higher ambient temperatures can actually reduce the risk of equipment overheating using automated monitoring with appropriate breaching thresholds, as local hotspots that may threaten key services can often go unnoticed if only a few temperature readings are available within the data centre.

Monitoring power and cost at server, OS and application level

A major problem with energy efficiency programmes is the lack of incentive from the end-users, departments or customers of a given organisation. Within the computing industry, it is rarely the case that end-users either have visibility or pay directly for their energy bill in relation to the IT services that are delivered to them. Therefore, it is very difficult to put in place a 'carrot and stick' approach that might affect their direct use of IT services in order to drive savings.

Sophisticated software that monitors power and cost at the server, OS or application level is able to provide both visibility and accountability in relation to energy consumption. Suitable charge-back models may also be deployed as an economic means to encourage change.

Defining IT efficiency metrics

Another difficulty faced by the IT industry relates to the definition of energy efficiency itself; not only with regard to the data centre as a whole, but also to the use of individual servers. Servers consume power in order to deliver a particular IT workload. However 'useful output' can often be difficult to define or measure, at least from a practical point of view.

In many common circumstances, a server or virtual machine is used solely to run a particular application and a suitable metric can be measured using protocols such as SNMP or WMI, or suitable application APIs.

Examples include:

- The number of emails sent or received by a Microsoft Outlook server.
- The number of database transactions by a MySQL server.
- The number of web transactions for a web server.

If the power used by the server or virtual machine is known or can be approximated, useful metrics such as transactions or watt can be used to define efficiency.

When suitable application metrics cannot be defined or are too complex to put into practice, simpler operating system metrics can be derived from combinations of CPU utilisation, disk I/O and network traffic.

CPU usage should be normalised by benchmark figures which determine the relative processing capabilities of different CPUs; this requires associated software to combine asset data with measurement data, so true efficiency metrics can then be calculated by using actual or approximate power for the associated servers and virtual machines.

Identifying under-used, faulty or inefficient IT equipment

While the industry has predominantly focussed on energy efficiency for cooling data centres, the largest potential savings actually come from the rationalisation of IT systems. In 2008, McKinsey estimated that average server utilisation was as low as 6%, since rising to 6%-12% based on anecdotal evidence from customers. A more recent paper by the Natural Resources Defense Council (NRDC) states that average server utilization remained static at 12 to 18 percent between 2006 and 2012.

Much of this stems from over-provisioning -where deployed servers are more powerful than necessary, or where no historical usage data was available to make an informed choice. In most cases, data centres have been designed for an assumed peak load with a very significant contingency factor, but in actuality they rarely approach a moderate percentage of this peak load. In most cases, server resources are little-used outside of peak hours, but they continue to consume significant amounts of electricity.

Evidence implies that a moderate proportion of the servers in a typical data centre are doing no useful work at all – the Green Grid estimated this to be as high as 10% in many moderate-size data centres following a customer survey. The NDFC report states that up to 30% of servers are 'comatose' – they are no longer needed because projects have ended or business processes changed, but are still plugged in and consuming electricity.

The identification of un-used, under-used, faulty or inefficient IT equipment is straightforward once a suitable monitoring system is in place. With the use of appropriate efficiency metrics, IT utilisation can be compared across systems, end-users, applications, computer architectures and even manufacturers. The peak, average and off-peak requirements can be accurately sized.

IT services for some back-office functions, such as payroll and invoicing, may also be rescheduled to take place out of peak hours, thereby reducing peak requirements. By planning IT rationalisation programmes and minimising procurement costs, data centres can thereby reduce energy usage, capital expenditure and operational expenditure.

Switching off or capping power for equipment when they are not in use

While it is clear that most data centres have some servers that are not used overnight, at weekends or in holidays, it is exceptionally rare for data centres to consider turning the equipment off during these periods. This is often due to the perceived risk of the servers not restarting properly, even though modern servers are designed to be power-cycled thousands of time during their lifespans.

Most modern Intel-based servers are also equipped with a power capping capability that can be controlled using out-of-band management networks. Software can then be used to set the maximum power drawn by these servers during off-peak times.

Power capping techniques are not only deployed as an alternative to switching off unused servers, but can also be used for services that require fewer computing resources during off-peak times – examples include email services, web-services and business-to-business services.

The implementation of power cycling and power capping strategies not only requires a good understanding of historical system usage (from operating system or application monitoring), but also an understanding of which computing resources are providing which resources. Such strategies can be implemented using semi-automated scheduling, or fully automated control systems.

Managing assets to understand OPEX costs against CAPEX costs

Most data centres contain servers which are three to five years old, with processing capabilities that may be an order of magnitude worse than a modern server. While the use of IT efficiency metrics would allow managers to identify the effectiveness of all of their computing estate, it is possible to define a number of useful metrics that combine asset information with historical power and usage information.

One such metric might take the net value of a server (the portion of the original asset cost that has yet to be depreciated) plus the estimated cost of power for that server over a given period (based on historical usage) and divide this by a server benchmark figure. This would correspond to an average cost/transaction metric for the period in question.

Similar metrics might also be defined for auxiliary equipment such as air-conditioning units, and would allow senior management to make financial decisions as to when old equipment should be replaced, while at the same time maximising energy efficiency. Related metrics might allow managers to understand CO2 emissions relating to their server estate, including embedded carbon.

Driving highly efficient virtualisation strategies

Recent reports and surveys by Gartner and the virtualization software provider VMware indicate that between 50% and 75% of the world's server workloads have now been virtualised. However, while virtualisation has helped improve the problem of low server utilisation, other surveys suggest that overall server utilisation remains low (a NRDC report suggests between 12% and 18%), and has stagnated in recent years.

While the claims relating to the use of virtualisation might imply a higher figure for overall server utilisation, it is important to note that even if 75% of workloads were virtualised, these virtual workloads would reside on significantly fewer than 50% of the physical servers. For example, with a ratio of five virtual machines per physical server, over 60% of servers would remain un-virtualised; the higher the ratio, the greater the percentage of physical servers that remain.

This is not to say that the total power consumption would not have reduced significantly in situations where demand for IT services have remained constant, just that the overall efficiency tends to stagnate. In reality virtualisation takes place when demand for services is actually increasing substantially.

In regards to IT efficiency an important issue is that the growth in hardware performance and capacity continues to outpace the consolidation of workloads. Given the evidence that many managers do not have good visibility of underused server hardware within existing data centres, it is likely that the majority of virtualisation projects are undertaken without a good understanding of peak and average load requirements.

Since virtualisation projects go hand-in-hand with hardware refreshes, IT managers may end-up replacing a single workload on a three to four year old server, with five to ten virtual workloads on a new server. However if that new server is capable of five to ten times the number of transactions, then the overall usage efficiency may not have improved significantly, if at all.

Without comprehensive, historical data on system and application usage, as well as ongoing monitoring, it is unlikely that virtualisation projects will achieve their potential for energy and CO2 savings. The Virtual Machine ‘sprawl’ is as likely to be as common as ‘server sprawl’ is thought to be today.

Active Control of IT and supporting infrastructure

While reducing the PUE and identifying under-used and inefficient servers can both lead to non-trivial energy savings, there is no doubt that real-time optimisation techniques offer the best prospect for saving energy within the data centre industry.

Given that computing demand can fluctuate significantly over time, as well as across applications and end-users, the software needed to operate a fully autonomous data centre would need to manage both the IT workload and the cooling infrastructure so as to optimise power continuously.

The requisite Data Centre Infrastructure Management software would need to:

- Collect fine-grain environmental information.
- Interact with UPS’s and cooling systems.
- Manage power from main distribution boards to the rack level PDUs and individual servers.
- Monitor server health, operating systems and even application performance.
- Manage dynamic VM migration.
- Approximate power utilisation by individual VMs and end-users.
- Move VMs to a specific part of the data centre during low utilisation periods.
- Power down or cap the power usage of under-used servers.
- Adjust cooling accordingly.

Potential Energy Savings

Given the many opportunities for minimising energy usage, the potential for a typical data centre to reduce electricity use and CO2 emissions is very significant.

If the initial focus were to be on cooling efficiency, most data centres should be able to reduce their PUE significantly with no major changes to the existing cooling infrastructure by optimising existing cooling systems, removing hot-spots and raising ambient temperatures.

Under the assumption that the average PUE of a typical data centre could be reduced from 1.8 to 1.6 or 1.5, and in the absence of any optimisation to the IT systems, a moderate energy saving of 11% to 17% could be made.

With a total redesign of the data centre, data centres should be able to operate at a PUE of 1.2 or below, for example through the use of free air cooling or water-cooled rack technology; albeit at significant capital expense, this would result in a 33% reduction in electricity consumption.

If the PUE is then maintained while the IT systems are rationalised, the total additional energy saving is then equal to the IT saving multiplied by the PUE. In a worst case scenario, we believe that a 10-15% reduction in IT energy could typically be attained with little or no capital expense. With more dramatic server consolidation, virtualisation and automation, a 50% reduction or higher should be possible.

The total potential energy savings based on a data centre with an original PUE of 1.8 are detailed below:

Achieved PUE	% Energy Saving (Cooling & Infrastructure)	Total energy saving following additional 10% reduction in IT Energy	Total energy saving following additional 15% reduction in IT Energy	Total energy saving following additional 50% reduction in IT Energy
1.6	11%	20%	24%	56%
1.5	17%	25%	29%	58%
1.2	33%	40%	43%	67%

Table 1: Potential Energy Savings for a Data Centre with a current PUE of 1.8

Perceived barriers to the adoption of energy-saving software in the Data Centre Industry

There are a number of perceived barriers to the adoption of DCIM products within the industry. A 2013 451 Group survey of 1000 data centre operators concluded that the following are the most important perceived barriers, order by importance:

- ***Cost is too high (>60%)***

  In reality, this relates not only to the cost of the DCIM software, but possibly more importantly the cost of putting in suitable monitoring hardware for power and environmental measurement. Potential customers do not seem to believe the potential for energy savings that are possible without significant capital investment, and therefore the ROI over a 1-3 year period. The best case energy (and OPEX) savings do rely on a much more significant investment in data centre infrastructure, virtualisation software and new servers, and while the potential ROI is very significant indeed, the willingness to invest in large-scale CAPEX projects is limited.

- ***Systems are not easy to integrate with existing equipment/software (30%)***

  Since many DCIM vendors, including Concurrent Thinking, support a large number of standard protocols in the data centre, the ability to integrate with existing hardware systems should in principle not pose an issue. However, most data centres would need to invest in hardware instrumentation, and some older systems may have hardware instrumentation that is unreliable – for example, a common problem with some managed power distribution units. Integration with existing software can indeed be problematic, partly because many building management systems are exclusive, and few have open interfaces. Some data centre owners have existing software that they have either developed in house, or bought from multiple suppliers, and there is reluctance in retiring these in favour of a common management platform.

- ***Our existing methods and tools already provide a sufficient solution (>25%)***

  This may be due to the potential benefit and breadth of DCIM, as few data centres other than the largest internet providers have efficient systems in place to monitor and manage their data centre. For larger data centres with existing systems in place, the risk of moving from such a system to a new software solution may constitute a significant barrier. Finally, we note that few of the respondents to the 451 survey were likely to be smaller data operators, many of which have no tools at all to monitor their data centres.

- ***Difficulty populating IT asset databases (>25%)***

  Many data centres have very limited systems for tracking assets (for example spreadsheets), and many have no such systems in place. Collecting asset information and then populating and updating a DCIM database can be time consuming, but this merely represents good practise.

- ***Systems provide insufficient benefits (>15%)***

  This relates the perceived ROI discussed above.

- ***Systems are too complex / unsuited to our needs (>15%)***

  This may be a function of the complexity of particular DCIM tools, or the difficulties installing these, as discussed above.

- ***Systems are not sufficiently proven (>5%)***

  With the relatively low take-up of DCIM solutions in the market, this is perhaps inevitable.

- ***Systems do not scale to meet our requirements (>5%)***

  Some of the larger data centre owners manage thousands of racks, and many DCIM vendors (especially those that have traditionally worked in the data centre facilities management space) did not design their systems to scale, especially given that monitoring IT devices requires two-three orders of magnitude more data points than 'facilities' equipment. Concurrent Thinking's products have been architected and tested to scale, although this has yet to be test 'in anger' at a large data centre.

The 451 Group concludes that potential customers may not budget appropriately and probably suffer "sticker shock". They also believe that customers find it difficult to compare prices because suppliers calculate the pricing in varying ways, and that there is also no "standard" set of DCIM features. Customers may also mistrust suppliers because of hidden extras and implementation costs.

From a supplier perspective, the other most significant barriers are the following:

- ***Split incentives and budgets***

  Most corporate data centres have multiple teams that look after different parts of the data centre, including: building management, facilities management, server management and IT/application management teams. In most cases, total energy utilisation is only known to the building and facilities management team(s), and it is rarely split into different cost centres. IT teams generally have no visibility of energy costs and are totally focussed on risk mitigation, at the expense of all other factors. Since DICM spans all of the above teams, there is often resistance to sharing management information, and there is rarely a budget for tools that span all of them.

- ***Long sales and implementation cycles***

  The sales cycle for DCIM seems to be between 6-18 months, and more often nearer the latter rather than the former, even for those companies that have a clear desire to acquire such solutions. Change management procedures within the data centre, especially IT change management procedures associated with networking and security, can cause very significant delays in the purchasing and subsequent implementation of DCIM projects. Availability of the data centre staff that will oversee and enable the project is also a key issue.

- ***Project priorities***

  While DCIM is becoming increasingly seen as a technology that data centres would like to and should adopt, it is often given a low priority compared with other data centre projects. DCIM budgets can also be diverted at the last minute due to unexpected capital projects: for example replacing failing air-conditioning equipment, or updating power distribution.

- ***Perceived risks***

  The implementation of many of the more advanced DCIM features, especially those can often lead to the largest energy savings, are often viewed as being too high risk. Data centre operators tend to take a very cautious approach to change rather than embarking on radical overhauls of their infrastructure. They are mindful of potentially variable and unpredictable workloads, and concerned about potential service level agreement (SLA) violations because performance could be affected by resource contention as server utilisation increases with virtualisation. They are also very conservative with regard to software that might help them automate procedures across both facilities and IT systems.

- ***Misaligned incentives between colocation data centres and 'cloud' providers***

  An increasing number of companies are now using colocation data centres to host their IT equipment. Historically colocation data centres have charged their customers solely based on space and peak power requirements; only a small percentage currently charge their customers for the actual power used, which results in little incentive for the latter to optimise their IT systems.

  Furthermore, if the data centre infrastructure is managed by the colocation company, and the IT systems by their customers, it is difficult for DCIM providers to offer their complete solution to either, and so the potential to maximise energy efficiency is lost. This problem will be mitigated by the use of the Colocation Portal technology that Concurrent Thinking have prototyped in this project, as this allows colocation data centres to provide detailed information about power consumption and environmental conditions to their customers, while allowing their customers to manage their IT systems within the same framework.

Data Centre Infrastructure Management and its relation to UK Energy saving incentives

The Government and the Department of Energy and Climate Change have proposed a number of schemes that incentivise industry to reduce energy use and carbon emissions.

These schemes or proposals include the following: Climate Change Agreements (CCAs), which reduce liabilities under the Climate Change Levy (CCL) or Carbon Reduction Commitment (CRC); Enhanced Capital Allowances (ECA), a tax rebate for capital purchases of classes of equipment that lead to reductions in energy consumption; and Electricity Demand Reduction (EDR).

Climate Change Levy (CCL) and Carbon Reduction Commitment (CRC)

From July 1st, colocation data centre providers with an electricity supply of at least 200kW have been able to take advantage of the Climate Change Agreement (CCA) that has been negotiated on their behalf by TechUK. CCAs allow companies to obtain tax reductions or exclusions from carbon taxes if they reach energy efficiency targets, and so reduce the impact of the Climate Change Levy (CCL) or Carbon Reduction Commitment (CRC).

According to TechUK, providers who pay the CCL and CRC will benefit by 1.35p per kWh of electricity, or about £24,000 a year for a 200kW data centre, and £120,000 a year for a 1MW data centre.

In order for colocation providers to benefit from this scheme, they effectively need to plan for and demonstrate a 15% improvement in PUE by 2020. This inevitably requires participating sites to collect auditable PUE data, although they are also required to regulate temperature and humidity within their facilities.

This represents a clear opportunity for DCIM vendors such as Concurrent Thinking as the government is incentivising the industry to build better energy and environmental monitoring into their day-to-day operations.

However while co-location providers operate data centres, they have little control of the IT equipment housed within them. They cool the data centre, but have minimal input in how much waste heat each of their customers will produce, other than by regulating the electricity that is available to each. Since many co-location providers' partition space and power using open cages, neighbouring cages could potentially generate more or less heat depending on how dense and efficient the IT infrastructure that is housed within each cage is.

Co-location companies who sign up to the CCA will inevitably spend time and money improving the effectiveness of the cooling systems within their data centre, which is ultimately good for them, good for the environment, and good for the industry.

However, if some or all of their customers decide to optimise their IT infrastructure, and the co-location provider doesn't respond to this, or cannot respond to it based upon its inability to target cooling for customers who need it more than others, then the PUE of the data centre may actually increase. If the saving that colocation customers make to the total IT load outstrips the relative saving to the cooling loads, then the PUE will increase. The data centre will make much more effective use of power than previously, CO2emissions will be lower but the colocation provider will lose the benefit of the CCA. This is clearly an unexpected repercussion of this relatively simplistic incentive.

Such changes would further benefit DCIM vendors such as Concurrent Thinking; the industry would not only need to have high level energy and environmental monitoring, but will also need to monitor fine-grain information, and potentially control data centre facilities to adapt to changing IT loads. It appears more incentives will be in place to ensure that metrics that take into account both IT efficiency and cooling efficiency can be generated by data centre providers.

Electricity Demand Reduction (EDR)

The original EDR consultation from DECC has now resulted in specific provisions within the 2013 Energy Act. These allow for financial incentives to be used alongside the CCL and CRC. EDR aims to tackle climate change alongside the UK's impending 'energy gap', by encouraging industry to consume less, or consume less during periods when others might want to consume.

EDR calls for government incentives to be made available for substantial and demonstrative improvements in energy efficiency. The caveat is that the incentives should only apply if the improvements stand up to an 'additionality' test: they should not apply to savings in energy that would have come about without the incentives. For example, if a colocation company simply had fewer customers, or if an obsolescent data centre was being wound down over time due.

Unfortunately most data centres cannot actually prove that they are efficient. Nor can they prove that any potential savings are 'additional'. The issue is that the required proof would ultimately involve knowledge of the 'useful output' of the data centre, which has to be a function of the computation being performed in the data centre. This is not the case for many other industries, where 'useful output' might be tonnes of sheet metal, litres of chemical compounds, or the number of car engines produced, and efficiency is then simply defined as 'useful output' divided by the total power consumed.

In order for the data centre industry to persuade the Government that these financial incentives should apply to it, it would presumably have to propose a pragmatic and simple measure of end-to-end IT utilisation. Based on the experience of the CCA agreement for data centres, the measure itself may not need to be perfect; some approximations of MIPs/kW, IOPs/kW and/or Mbytes/W might suffice.

Energy Demand Side Reduction

Demand Side Reduction (DSR) is a mechanism by which large energy consumers can voluntarily reduce the amount of electricity that they use during peak times (typically between 1600hrs and 2000hrs on winter weekdays). The UK National Grid is actively seeking companies that can commit to such reductions on a voluntary basis, in return for a payment. The scheme will ultimately help keep energy system costs down for consumers by avoiding the need to build additional power stations. Companies who participate in the scheme not only benefit from a direct payment, but also reduce their bills because of the higher cost of energy at these times.

DSR has been successfully demonstrated in the commercial sector by companies such as Kiwi Power, which provides software that interfaces with customers' building management systems to reduce the energy used by 'non-essential' systems on demand. A reasonable example of this might be the air conditioning systems used by a chain of hotels. Most hotel customers wouldn't notice or complain if the temperature of a hotel lobby increased or decreased by a few degrees over a short period of time. If this scenario was enacted across a modest number of hotels, using software that orchestrated the 'how and when', then a significant amount of energy could be saved without having a noticeable impact on the hotels' customers.

The potential issue with DSR in the data centre industry relates to what might be considered a 'non-essential' service, as data centres as a whole are often considered as mission-critical to the organisations that use them. DSR proponents have nonetheless mooted the possibility of reducing electricity during times of peak demand by using back-up generators or UPS systems to provide power to the data centre. However, this simple concept, while attractive at first glance, may be fool-hardy given that those generators and UPS systems were specified so as to maintain critical services in the event of a power failure. Reducing the time period over which they provide back-up cover is presumably not an option if the systems were properly specified in the first place.

In reality, if DSR is to be of widespread use in data centres, owners will need to manage the IT services that are provided by their data centres in a much more pro-active and efficient manner. They will need to understand what IT services are running when, and be prepared to schedule such services differently while maintaining the same overall computational throughput.

For example, if historic utilisation figures show that there is significant underused computing capacity overnight or at weekends, certain tasks such as payroll or commercial reporting tasks could be rescheduled from peak times to these less busy times, leaving more flexibility for back-up energy systems to be run down while maintaining the required cover for critical services during peak times. Air cooling systems could also be focussed on maintaining these critical services, while other parts of the data centre could be allowed to increase in temperature.

So while DSR is both possible and potentially very worthwhile for data centres, it requires a joined-up approach to managing both facilities and IT systems, as well as a clear understanding of which end-users or applications require which resources, when.

Bibliography

- Digital Realty Trust (2013) North American Campus Survey Results
- Jonathan Koomey (2011) Growth in Data Center electricity use 2005 to 2010. Oakland, CA: Analytics Press
- Forest et al, McKinsey and Company (2008) 'Data Centers – How to cut carbon emissions and costs'
- NDRC (2014) 'Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers'
- Gartner (2013) 'Virtualization Penetration rates in the enterprise, 2012' Cloud, IAAS & Virtualization Survey
- VMware (2011) 'VMware Business and Financial Benefits of Virtualization'

Contact us

To find out more about Concurrent COMMAND or to request a demo, contact us on sales@concurrent-thinking.com.