ORIENTATION GUIDE

Energy Aware Planning and Decision Making

Five imperatives for IT leaders

Dale Vile and Tony Lock

Freeform Dynamics Ltd

May 2012

This paper was authored on an independent basis by the analyst team at Freeform Dynamics Ltd, with sponsorship from CA Technologies. The content is based on the industry intelligence gathered and analysed by the authors and does not necessarily reflect the views of the sponsor.

Introduction

The question of energy availability and management is often considered to be a core component of the so-called 'green' discussion. This is a mixed blessing. For the minority of organisations with a well thought out and properly resourced 'green agenda', the drive for energy efficiency is often the most tangible part of the business case for taking action, as the link to cost savings is clear. For others, wrapping the question of energy management up with what is sometimes seen as a 'cause', or a problem too nebulous to pin down and act upon, means it never gets properly addressed.

When considering energy consumption, we then have to consider 'organisational politics'. Most managers and executives today are not directly accountable for the energy used in their part of the business, and the last thing anyone wants is something else to worry about managing. If someone does have to step up and take responsibility, then the secret wish is "let that someone be someone else!" When the question of IT-related energy use is raised, the CIO can then easily end up in the hot seat, even though power consumed by systems is largely determined by business activity.

If this description of corporate reality strikes a chord with you, and you are finding it difficult to get off the ground or make progress with energy-related initiatives within IT, then this paper is for you.

Our aim is to illustrate that taking action on energy need not be onerous or unnatural, particularly in the context of the data centre where the emerging discipline of 'data centre infrastructure management' (DCIM) provides solid foundation for driving energy-related improvements. As we shall see, a lot can be achieved by taking a more joined up 'DCIM style' approach to management, which basically means considering power consumption, cooling, capacity, energy supply and other key aspects of data centre operations together, all in the context of IT service delivery requirements.

Before getting into that, though, let's recap on why doing anything at all in this area is important.

The energy management call to action

Businesses rely on IT, and IT relies on energy – often quite a bit of it in the case of corporate data centres. The three obvious consequences of this are:

- Energy consumption via the IT infrastructure is a significant cost, and therefore a potential constraint on IT-related activity when pressure to manage that cost is applied.
- Energy availability, i.e. being able to source enough energy to power the IT infrastructure and be confident in its supply, is fundamental to both growth and business continuity.
- Energy-related regulation will almost certainly have a direct or indirect impact on IT.

With this in mind, the proactive management of energy consumption is going to become increasingly more critical. Indeed, while it is often overlooked, the need to take more control is already pretty important given global trends and developments. Consider, for example, the depletion of finite fossil fuel reserves, the playing out of political and conflict scenarios between nations, and the physical challenges delivering power to densely populated areas in many countries. We can expect trends and events of this kind to lead to higher energy prices, unpredictable price fluctuations, possible caps on availability, periodic interruption of supply, and legislation from governments to act as a 'stick' for beating 'irresponsible' businesses (Figure 1).

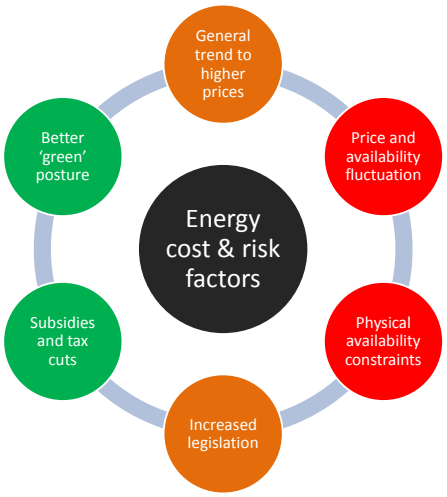

Figure 1: Energy cost and risk factors

However, opportunities will increasingly emerge to take advantage of subsidies and tax breaks offered as incentives by power companies and governments to encourage responsible energy use, which we can think of as the 'carrot' to complement the legislative stick. It is also clear that for those with a broader environmental agenda, better energy management will

enhance their 'green posture', with benefits arising from the creation of a stronger positive reputation with both customers and stakeholders.

So, paying more attention to energy makes good business sense. But how does this translate to specifics, particularly in relation to IT?

Factors affecting energy consumption

In order to act on some of the risks, costs and opportunities we have been discussing, it is first necessary to be clear about why energy management is so often lacking in the context of IT. While you may not have immediate problems and shortcomings in all of the areas identified in Table 1, we would be surprised if you didn't recognise at least some of them.

Problem	Comment
Lack of visibility	As the old adage goes: "You can't manage what you don't measure", so if there is limited or no visibility of how much energy is being consumed by different parts of the IT and facilities infrastructure, then it's hard to know how to prioritise and focus efforts to improve things, or to know when results have been achieved or objectives have been met.
Lack of management ownership	Energy accounting is often based on allocation of costs to departments according to parameters like headcount or square-footage of office space. With no direct accountability or mandate to control, business managers have no incentive to work together with the IT function to optimise IT-related energy use.
Inherently inefficient IT equipment	IT assets accumulate over time, and you can often find several generations of kit coexisting in the data centre or computer room. Older servers, storage devices and network switches are generally more power hungry and harder to manage, so the more of them there are, the less power efficient the overall infrastructure.
Poorly utilised IT equipment	While virtualisation has been used to consolidate portions of x86 server estates, the truth is that the vast majority of servers and storage devices on the planet have not been virtualised. A great deal of equipment is therefore still sitting there consuming power continuously but doing little real work for most of the time.
Over-stretched capacity	After years of continually adding IT equipment into the infrastructure with little concern for space, power and cooling, many data centres are now operating at or near capacity. Provisioning new systems therefore is often now a challenge, creating delays responding to new business requests, along with conflicts and 'max out' events that cause unexpected service interruption.
Poorly utilised applications	Organisations tend to have a continual stream of project work to deliver new applications, but pay much less attention to identifying candidates for application consolidation, rationalisation or end of life. Many old applications therefore exist that deliver little value but still consume IT resources, and therefore power.
Parochial attitude to IT efficiency	When local interest overrides the greater good because of politics, convention or lack of awareness, opportunities to rationalise, consolidate and share can meet significant resistance. Problems with parochialism also occur when IT is unwilling to spend a little more internally to enable bigger savings across the business.

Table 1: Common causes of poor (or non-existent) IT-related energy management

This is not an exhaustive list, but what we see here are the most common causes of energy-related risk and inefficiency we come across in our research activities at Freeform Dynamics.

Now shrewd readers will have spotted something very important here. While the focus of this paper is on energy management, the factors listed above largely coincide with the underlying causes of inefficiency and ineffectiveness within IT in general. To put it another way, when looking at energy management, we're mostly going to be considering the same set of problems and principles that we would when formulating any broader IT delivery improvement initiative. And by the same token, we would also be looking at a similar set of enabling approaches and solutions.

Going with the flow, rather than against it

At this point, it is useful to review what's been keeping many IT departments busy over the past few years, along with the kind of things on the agenda to boost IT efficiency, create a flexible and responsive environment, and maintain or improve overall service levels to the business.

These include: revisiting IT governance with a switch in emphasis from systems to services, virtualising and modernising the infrastructure, exploiting the benefits of shared infrastructure and services, and tackling IT management in a more joined up manner. More recently, the maturing of ideas, technology and services around cloud computing has been encouraging a more dynamic approach to IT delivery and management, whether through flexible third party hosting, or the adoption of 'private cloud' architectures to optimise the use of internal IT assets and resources.

The good news is that as all of this unfolds, organisations will quite naturally become more energy efficient, whether they have this as an explicit goal or not. With more precise understanding of the factors that make the most difference, however, a more energy aware approach to decision-making and investment can be made. Activity can then be tuned and more can be achieved without the need for additional investment beyond that already scoped. And obviously the sooner you start, the less likely it is that the organisation will be caught out by fluctuating or escalating energy costs, emerging legislation on the use of power or other energy-related risks.

So how does this translate to an action plan for those wishing to develop or optimise their approach to proactive energy management?

Five key imperatives for effective energy management

Ironically, one of the challenges with energy management is that so many options for potentially driving improvement exist that it can sometimes be very hard to prioritise activity. After all, if you are going to spend time and effort in this area, you want to be sure that you achieve the most impact, and leverage off existing commitments and activities as much as possible.

In order to help with this, we recommend focusing on five imperatives in relation to IT (Figure 2).

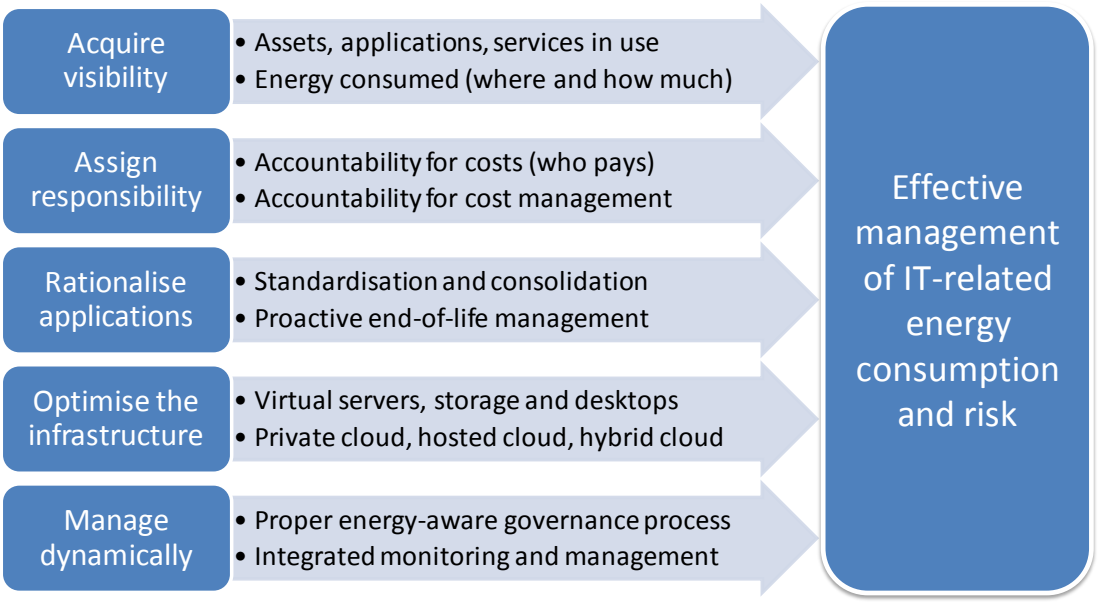

Figure 2: Five imperatives for effective energy management in relation to IT

Based on our previous discussion of the common causes of inefficiency and risk, the imperatives presented here should be largely self-explanatory to those with a good working knowledge of IT management. However, it is worth calling out some specifics in relation to each (particularly the first two), so let's walk through them:

1. Acquire visibility

Let's be blunt: you're going to get nowhere with attempts to boost efficiency and reduce risk in IT if you don't have a good handle on what's in your IT, and facilities, infrastructures and how it is used. Adequate asset and configuration management is therefore an essential pre-requisite for any attempt at proactive energy management to succeed. Asset discovery solutions and configuration management databases (CMDBs) are the kinds of tools that can help here at a systems level, with application and service portfolio management at the next level up (i.e. the IT-business delivery interface).

Building on this, you can then start to use DCIM tools and techniques to overlay energy consumption information and form a body of intelligence and insights that will facilitate energy aware decision-making. At the very basic level, especially in simpler or smaller-scale environments, a lot can be achieved by just recording power rating information against key assets so you can form a view of which parts of the infrastructure are hungriest in terms of power. Beyond this, monitoring of actual energy consumption allows understanding to be taken to the next level. Today, with the right monitoring technology, it is possible to measure power consumption at pretty much any level of granularity, from the overall data centre, through individual racks, to specific items of equipment. Similarly, tools and techniques exist under the DCIM umbrella to measure and analyse the thermal characteristics of the data centre environment at a fine grain level to identify 'hotspots' before they pose significant risk.

What's right for your organisation will depend on how complex it is and how much detail is required to enable energy related decisions to be made effectively (more of that in a minute).

2. Assign responsibility

Many would argue that the assignment of responsibility is actually the place to start, and that's certainly valid in principle, at least in terms of executive level ownership. We have listed it second, however, for two reasons. The first is because the ownership and responsibility discussion at the next level down is often dependent on at least a basic level of visibility – e.g. which applications or services (and by implication underlying assets) are used by which parts of the business and how much. The second is because you may choose to fold energy cost and risk into the broader discussions concerned with IT delivery and accounting in general.

On this last point, there's a lot going on in IT which is putting pressure on traditional ways of doing things. The prevailing trends mean that IT needs are increasingly going to be met through a combination of external services and internal infrastructure that is shared rather than dedicated. This in turn means an inevitable switch in focus from the cost of running 'systems' to the cost associated with the delivery or consumption of 'services'.

Whether this should manifest itself in the form of cost allocation or more explicit chargeback is an ongoing debate. Either way, it probably makes most sense to regard energy consumption as a component of overall service costs, but one that is visible. If energy-specific objectives are then set, the notional power rating of supporting services can be taken into account by those responsible for optimising business processes and functions. An understanding of the power dependency of key services can also be used when making decisions around business risk.

3. Rationalise applications

When multiple facilities, software systems or physical assets exist that provide the same or very similar capability, there is a clear opportunity to consolidate. It is beyond the scope of this paper to get into the details of what's involved in consolidating data centres, ERP instances, email systems, etc; suffice it to say that such activities can often lead to a significant reduction in energy consumption, as well as other costs, if approached in the right manner.

Beyond consolidation, it's also often possible to simply shut down certain systems. This might sound quite dramatic, but the reality is that applications are frequently created to serve a transient need so come to a natural end of life, even though this is often overlooked. Old applications that are today only used out of habit rather than real need are also candidates for termination, especially if they are sitting on old equipment that is both inefficient and over-specified because it was originally put into place to deal with high levels of activity that have since dwindled.

This kind of rationalisation and consolidation is something IT departments should be doing anyway. In the context of energy management, the only thing you would do differently is include power consumption as a criterion for efficiency assessment, and/or increase its weighting as necessary.

4. Optimise the infrastructure

Modern server and storage virtualisation techniques have already gained popularity because of the reduction in cost and overhead they enable. Virtualising the IT environment is also good news from an energy perspective. Workload consolidation means fewer servers need to be powered to support the same range of activity. Storage virtualisation means more effective use of the capacity available in the infrastructure, again reducing power consumption through better utilisation. Today storage is often the fastest growing part of the IT infrastructure, but can be overlooked.

In addition to the above benefits, virtualising the IT infrastructure means that new application requirements can often be met without the need for new equipment, so the efficiency gains are continuous in nature. Such benefits are further enhanced as organisations embrace the concept of private cloud, which allows resources to be allocated and reallocated in a highly flexible way as demand fluctuates, with a further boost to utilisation, which in turn optimises energy use even more.

Quite apart from virtualisation and cloud, simply modernising server and storage hardware estates can lead to significant power efficiency gains. The latest equipment tends to be both more power-efficient, and more manageable from an energy perspective, e.g. allowing power consumption to be ramped up or down depending on activity levels. More efficient equipment also means less heat, and when you put this together with the fact that server and storage devices today can generally run safely at higher temperatures, much less energy is required for cooling.

When looking at hardware and configuration options, there is a need to make sure that workloads are placed on the most efficient and effective platform. Beyond the workload characteristics and hardware/software dependencies, it is important to consider the energy profile of candidate platforms, bearing in mind the power required to cool the equipment as well as the electricity it will consume directly. Sometimes it will make sense to use a large Unix box, sometimes a traditional cluster or perhaps even a private cloud. And if you have a mainframe, it is important not to overlook this, as today's mainframe architectures can run a variety of workloads, including Web and Linux, and will often represent the most power efficient option.

Evaluation of platform and placement options using various forms of 'what if' analysis is an area in which DCIM tools can be particularly helpful.

5. Manage dynamically

The reality is that in any IT environment, demands are frequently changing, and it is a failure to keep up with this change that often leads to redundancy and waste. The answer is to put in place processes and tools that allow everything to be managed holistically from top to bottom on a continuous basis (Figure 3).

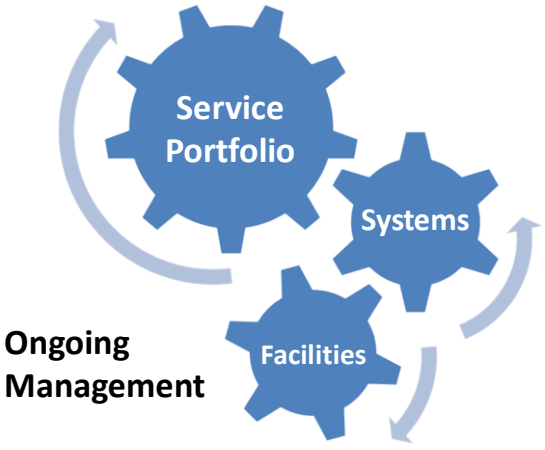

Figure 3: Holistic approach to ongoing management

At the top of the stack, this translates to a need for an effective and ongoing IT governance process, and the emerging discipline of 'service portfolio management' is a key enabler here. The main difference between this and traditional project portfolio management is that attention is paid to existing systems and services as well as those undergoing development or change. The trick is to continually challenge whether what's in place is optimal for meeting current needs, and it is through this that opportunities to consolidate, migrate, terminate, re-platform, or otherwise optimise application delivery are identified.

From an energy management perspective, the requirement is to include the consideration of power and cooling efficiency and energy related risk when reviewing existing services or assessing investment propositions for new ones. These new criteria sit alongside the usual cost and risk parameters, and the end result is an 'energy aware' planning and decision making process.

Lower down at the systems and facilities levels is where alternative sourcing and implementation options should be considered, and the freedom to make correct choices (and thus optimise for energy efficiency) is heavily dependent on the appropriate management capability being in place. It is important to be able to facilitate provisioning, monitoring, troubleshooting and support across different internal and external domains. This may include cloud based scenarios in which on-premise and hosted resources work together (the so called 'hybrid cloud' model).

Lastly, end-to-end monitoring and management of the systems activity underpinning services, using where appropriate some of the DCIM tools and techniques we have mentioned, is a pre-requisite for generating visibility, as well as assuring end-to-end service quality.

Final thoughts

In this paper, we have focussed on energy related efficiency and risk in relation to IT, and walked through some of the associated imperatives. It is important to remember, however, that in many environments, the IT element is relatively small when compared to the overall level of energy consumed by the organisation as a whole. While we urge action in the areas we have highlighted, IT must not think or act too parochially, as there will be occasions when consuming more power in IT will enable greater savings elsewhere. With this in mind, the key is to take a balanced and collaborative approach.

About Freeform Dynamics

Freeform Dynamics is a research and analysis firm. We track and report on the business impact of developments in the IT and communications sectors.

As part of this, we use an innovative research methodology to gather feedback directly from those involved in IT strategy, planning, procurement and implementation. Our output is therefore grounded in real-world practicality for use by mainstream IT professionals.

For further information or to subscribe to the Freeform Dynamics free research service, please visit www.freeformdynamics.com or contact us via info@freeformdynamics.com.

About CA Technologies

CA Technologies is an IT management software and solutions company with a deep expertise across all IT environments—from mainframe and distributed, to virtual and cloud. Solutions from CA Technologies enable customers to manage and secure IT environments and deliver more flexible IT services to the business faster.

Solutions from CA Technologies are the critical foundation for helping customers gain deep insight into and exceptional control over their complex, mixed IT environments. Our solutions enable customers to support changing business models at the right time, the right place, and from the best sources, and flexibly deliver IT services in the manner that best serves business needs.

Our solutions and services are the foundational elements for the deep insight and exceptional control necessary to make the promise of flexible IT and agile business a reality.

To learn firsthand how CA Technologies helps our customers succeed, please go to ca.com/success.

Terms of Use