


Innovation in your data centre

A White Paper



Date : February 2008
Reference : GB/02/014

Contents

Contents	2
Innovation in the data centre	2
The role of “innovation” in the data centre	3
Different innovation styles suit different situations	3
The evolution of the data centre	4
Smaller/Quicker/Cheaper	4
Automation: Living up to decades old promises	4
Green IT – It's as much about the bottom line as saving the planet	5
What is “Green IT”?	5
Today; Green is about cost	5
Data centre heat and power is just the start	5
A holistic view of Green	6
The infrastructure layer	6
You have to consider the full life-cycle at the infrastructure layer	7
The process layer	7
People	8
Infrastructure 2.0	8

Innovation in the data centre

We live and work in a fast-moving and dynamic world, in which we're called upon to support a level of change in business processes that we've never seen before.

In order to support these demands; what role does the data centre play in delivering “business innovation”?, and what changes can we expect to see in our data centres over the coming few years?

In this paper, Gary Barnett discusses the role that innovation can play in the data centre, the potential impact of virtualization, and how “greening your data centre” is as much about delivering efficiency and effectiveness as it is about saving the planet.

The role of “innovation” in the data centre

The term “innovation” is used far too often within the IT industry, and end-user organisations are right to be sceptical about the claims that many vendors make about their ability to deliver “innovation”.

To begin with, it's important to look at “Innovation” in your terms – rather than the vendor's. Innovation is only meaningful to you if it results in some benefit to the business and this is where some vendors come unstuck; they're happy to talk about their innovations as they apply to their products, but not so ready to talk about how that “cool new stuff” translates into benefits for you.

The idea of innovation is even more complicated when talking about technology infrastructure. Innovation is associated with attributes which often don't belong inside a data centre - “Exciting”, “Creative”, “Adventurous” aren't terms that automatically inspire confidence when you're talking about the core network, your disaster recovery plan or your approach to service level agreements. Relatively few COOs would be pleased to hear that you've just implemented an “exciting new disaster recovery process”, or that you've found an “adventurous solution to meeting your service level agreement”.

We define innovation as follows:
“Innovation is the application of creativity; resulting in improvements in service, process or product that deliver tangible and measurable business outcomes”

Different innovation styles suit different situations

The first point to make however, is that there are a range of innovation styles – at one end of the spectrum lies “Radical Innovation” and at the other end is “Incremental Innovation”.

Radical innovation implies a major change, which goes hand in hand with significant potential benefits, and often equally significant risk. Incremental innovation is more about gradual, step-wise, improvements; smaller gains, but much lower risk.

In domains where risk is a major concern (data centre operations would certainly fall into this category) you would not rush ahead to implement dramatic changes. It would clearly make sense to bias your innovation style towards the incremental approach where risk can be effectively contained.

But, there may well be cases where more radical forms of innovation are called for and provided that you are confident that you can manage the risks associated with more aggressive change there's no reason why you shouldn't go for it.

Innovation in the data centre is about efficiency and effectiveness

While there is rarely a direct connection between improvements that you might make in your data centre and the “business outcomes” that the COO is looking for, innovation does play two important roles; the first, and more obvious one, relates to efficiency which, from the COO's perspective, translates into cost. The second relates to effectiveness – the ability of IT to respond to business demands. Two examples of efficiency initiatives include:

- **Server consolidation**

Using technologies like blade and virtualization to reduce the number of server machines – lowering hardware and server administration costs

- **Automation**

If you can automate some, many, or all of the core administrative functions (server back-up, base administration etc.) you can significantly reduce the manpower costs associated with server admin.

“Efficient IT” is IT that uses the least resources, and runs at the lowest cost.

But efficiency doesn't automatically equate to effectiveness – there's not much advantage in running an efficient IT shop if, as a result of all that efficiency, it is inflexible.

The only thing you can be certain of (apart from death and taxes) is that the requirements of the business are going to change. It's normally safe to predict that much of that change is going to come in the form of a requirement for more – more storage, more processing power, and more flexibility.

So ensuring that you have room to scale is essential if you want to avoid having some difficult conversations with the business in a few months time.

Effective IT is, usually, efficient – but it's more than just cost-efficient, it's also capable of supporting new business processes and the applications that support them.

Effective IT may not be a direct source of business innovation – but it has a huge role to play as an *enabler* for business innovation.

Innovation in the data centre is as much about enabling business innovation as it is about applying new approaches to delivering efficient IT services. You should choose to implement technologies that enable you to support the growing IT needs of the business for storage (which more or less doubles every 18 months) and for processing power which is growing at a similar rate. Key activities include:

- **Consolidation**
- **Automation**
- **Managed and virtualized storage**
- **Blade technologies**
- **Server Virtualization**

The evolution of the data centre

A number of factors are going to have an impact on data centres;

- **Smaller/quicker/cheaper**

The ongoing improvements in hardware design will continue to deliver significant improvements in price/performance

- **Automation**

The complexity of our systems, coupled with technologies like virtualization will mandate the use of sophisticated systems management tools

- **Green IT**

While for many, particularly in North America, the green IT agenda is driven by economic factors, elsewhere in the world environmental regulation is already planned, and data centre managers need to be prepared

- **Infrastructure 2.0**

The emergence of web-based computing and Software as a Service (SaaS) will drive changes in the way we provision our data centres, and how we secure them.

Smaller/quicker/cheaper

For years, many of the drivers for change in the data centre have been a result of the continued improvements that vendors have been able to make in server and chip design.

Chip designers have had to face a major challenge in the past two years; the demise of Moore's law. It became clear in the 1990's that it wouldn't be possible to continue to build faster processors with ever increasing numbers of transistors, not least because the gaps between transistors on microprocessors can now be measured in atoms. So the tried and tested trick of doubling everything (which includes doubling the amount of heat generated by the processor) no longer works.

The good news, however, is that you can expect vendors to continue to provide similar levels of price/performance improvement to those we've seen for the last three decades; but they'll have to deliver those improvements using different techniques.

The emergence of multi-core general purpose processors (of which IBM's Power4 processor was the first) along with a much more systems-oriented approach to CPU design will deliver a similar growth in CPU performance. But these improvements do come at the cost of slightly greater complexity when it comes to writing software, and existing applications may not get the full benefits of the added performance unless they are recompiled, and in some cases modified, for the new processor designs.

Improvements in processor design will continue to deliver price/performance improvements – but in order to gain the benefits of the improved performance, some software will have to be modified

Automation: Living up to decades old promises

For years, automation has been hailed as a way not only to dramatically lower the cost of running infrastructure, but to make it more reliable and scalable as well.

While the market for systems management tools has seen good growth over the past 10 years, it is still under-exploited by most organisations.

The growing complexity of the infrastructure that organisations have to support coupled with the growth of virtualization technology, means that systems management tools will become mandatory for most organisations.

Some organisations have already implemented the “lights-out data centre”, where all of the systems operate with next to no manual intervention. The number of organisations that follow this trend is set to grow sharply over the next few years.

Virtualization technology, once the sole preserve of the IBM mainframe, has entered the mainstream. Virtualization makes the challenge of server consolidation much easier by providing an environment that allows many virtual servers to run on the same hardware without running the risk of them clashing. Virtualization tools are becoming increasingly sophisticated, and the ability to move workloads from one machine to another without interrupting service will be present in mainstream products within the next couple of years.

Virtualization technology is now mainstream – over the next three years IT’s capabilities and sophistication will improve significantly. Combined with systems management technology, the ability to dynamically allocate and prioritize workloads will be available to medium-sized and smaller data centres.

But, it is possible to overstate the benefits of automation. The biggest cost saving should be in administration costs, but although automation and virtualization may lower the costs associated with running the actual hardware, it is unlikely to reduce the need for human administrators once it’s virtualized.

Green IT – it’s as much about the bottom line as saving the planet

Green IT has become a major marketing band-wagon for IT vendors; all of whom are jostling to position themselves as “greener than thou” in the market place.

What is “Green IT”?

The Green IT debate is currently dominated by the IT hardware manufacturers, and is primarily driven by energy costs rather than environmental concerns. The “Green” label many vendors attach to their products is largely driven by marketing.

Whether these vendors are simply jumping onto a green bandwagon or not, it is clear that they have a key role to play in helping drive down the growing energy costs associated with IT.

Today, Green is about cost

It’s important to note at this point that we’re not going to justify action on the basis of the benefit it will bring to humankind (although it is certainly true that every little counts). We’re promoting a green agenda on the basis of the business benefits you’ll derive from it. The good news is that there will be plenty of opportunities to lower costs while helping the environment at the same time.

Data centre heat and power is just the start

While we agree that looking at your data centre energy footprint is the right place to start, it’s important to remember that over half of IT’s energy bill comes from outside the data centre; i.e. from the network and distributed computers. Certainly, the key medium term priority is energy cost, but government regulation and market demand will grow in importance over the next five years (particularly in Europe), to the point that they will become the dominant factors in determining your green agenda.

While we’re surprised by the number of organisations that *don’t* measure data centre power consumption, the situation is far worse elsewhere. Few organisations can

come anywhere close to measuring the energy consumed by the PC's, routers, printers etc. that are spread across the organisation. And bear in mind that just as your servers need cooling, so do all of those PCs.

It's important to remember that IT has a contribution to make beyond the boundaries of your technology infrastructure; in developing your green agenda you have to take a broader view.

If you want to reduce the overall energy footprint and the resultant energy bill, you need to look at the data centre in its entirety.

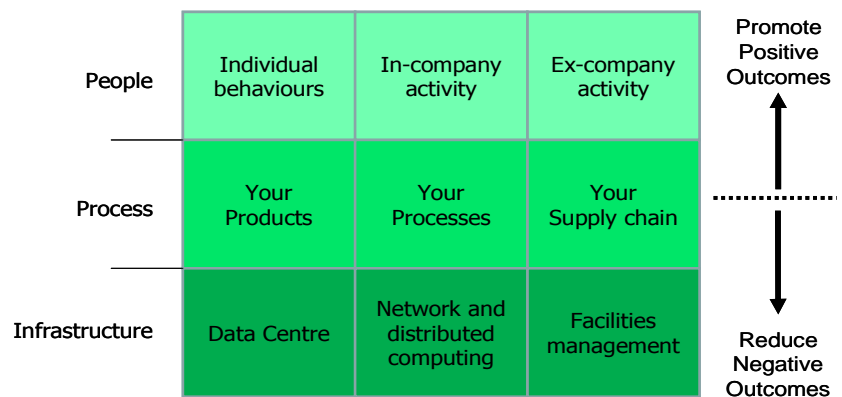
One of the biggest areas of energy consumption is cooling, for example. The majority of data centres are inefficiently cooled. Many are kept at a lower temperature than really required (a humidity of 45-55%, and a temperature range of 20-25 degrees C represent typical recommended ranges) or have cooling systems

A holistic view of Green

On the right is the Bathwick Green Matrix, which highlights the range of issues that need to be addressed in order to develop a comprehensive green agenda.

The model is divided into three layers, beginning with Infrastructure at the bottom, which deals with the physical things within your organisation that consume energy. The next layer – Process – addresses the fact that the way you do business has an impact on your environmental footprint. The top layer – People – deals with the different behaviours promoted by the company.

The Bathwick Group Green Matrix



© The Bathwick Group 2008

Broadly speaking, there are two types of initiative; those that can potentially deliver business value, and those where cost effectiveness is the primary benefit. Typically improvements made lower down in the matrix affect cost, while those in the upper half bring business value.

The infrastructure layer

Data centre

The data centre is currently the main focus of the leading hardware vendors – All of whom are marketing their technology, at least in part, as “energy efficient”.

There are a number of potential issues with this approach to marketing; making bold claims about the power consumption of a processor will lead to heart ache if the overall power consumption of the data centre isn't significantly improved.

that are simply not appropriate to the more modern, and more compact, computers that they are intended to cool.

There have been several recent case studies, from a number of vendors, where the re-design of a data centre's cooling infrastructure has brought about major energy savings.

There are other ways to improve the energy efficiency of the data centre which can also bring significant savings in total cost. Work-load consolidation, for example, can reduce hardware cost as well as reducing energy use.

Distributed computing environment

It is more than likely that the computing and network devices that live outside the data centre will consume significantly more energy overall than the data centre itself.

Desktop computers generate heat which has to be offset by the air conditioning system in the summer months; conversely they reduce the amount of heat that needs to be circulated during winter months.

The problem is that, by its nature, it is far more difficult to create significant energy savings in the distributed computing environment – but just as with the data centre, there are things you can do that will bring multiple benefits.

Thin-client technology, for example, uses significantly less power than PCs do. In some cases thin clients (those with an integral monitor) use less than a quarter of the energy of the PC they replace. Of course, this places a greater load on the data centre, but there are additional benefits to be gained (reliability, security, cost of maintenance and support, admin etc.) by moving to a centrally-served thin-client model for desktop computing.

Facilities

IT can also play a role in reducing the overall energy bill for the facilities you use as a whole. Effective environmental control (which might begin by simply ensuring that parts of the building that aren't in use aren't heated in the winter or cooled in the summer) can make a significant difference to overall energy costs.

New buildings are increasingly designed with energy efficiency in mind, while older buildings are difficult to modify in order to make them more energy efficient. But there are still things that can be done to reduce overall facilities energy costs; from simply automating the turning off of lights, to redirecting warmer air expelled from the data centre into the main A/C system in the winter (and ensuring that this doesn't happen in the summer).

You have to consider the full life-cycle at the infrastructure layer

It is important that you don't focus your attention exclusively on environmental profile of your technology while it is in operation. Every technology product has a full life-cycle from dust (where the raw materials are assembled to make the device) to rust (when the device is finally disposed of).

Increasingly, environmental legislation will require that you take an interest in the full life-cycle of the products that you use. The focus of the legislation, at least for the next few years, will be the end-of-life phase but it is clear that the original manufacturing process can have a significant impact on the cost of disposing of "retired" computer hardware.

If you don't already have a hardware disposal policy that takes into account environmental impact, you will need to develop one within the next couple of years.

The process layer

This layer contains, perhaps not surprisingly, the business processes and activities that are involved in the running of your business, the creation of the products or services that you sell, and the management of your supply chain.

As environmental impact becomes a more important factor in consumer choice you will increasingly be expected to be able to measure, and justify, the environmental impact of all of the activities that take place within this layer.

Your products

The focus here is efficiency on one hand, and innovation on the other. Technology can play a role both in delivering more efficient production processes or in supporting the development of new (and efficient) processes. How can you use technology to ensure that the products and services you provide have a lower environmental impact than your competitors' offerings?

Your processes

Within five years you will be asked to account for the environmental impact of your internal processes. In Europe, suppliers to the public sector are already being asked to provide environmental impact statements. We expect these to become mandatory for all major government sector contracts within the next five years.

You need to consider how your processes can be improved so that they are both more efficient in general terms (including their environmental impact) and more effective in terms of their flexibility, and support for change.

Your supply chain

As your customers increasingly demand that you account for the whole impact of the products and services you offer, you will have to make the same enquiries of your suppliers.

When you sell a product or service, you will be increasingly required to account for IT's environmental impact throughout its whole lifecycle.

People

This layer is primarily focussed on behavioural/cultural change. Again, technology has a key, albeit indirect, role to play in supporting this change.

Individual behaviours

What can you as an organisation do to ensure that the individual behaviours of your employees have a positive impact on the environment?

For example, company policies about home working not only saves significant cost in terms of premises, but also have the potential to reduce "commute miles" and consequently lower the overall carbon footprint of your organisation.

In-company activity

What can you do internally to foster the cultural change necessary to ensure that everyone within your organisation is motivated to be environmentally aware (and rewarded when they are).

Extra-company activity

Corporate and social responsibility is no longer for internet start-ups and hippy communes. Consumers and shareholders are increasingly interested in the contribution that companies make to their local communities. For companies investing in China a willingness to commit to supporting the development of the local community is already more or less mandatory if you want to gain government support for you initiative.

Green IT isn't just about making the world a better place; it makes sound commercial sense.

Start by looking at the infrastructure layer and initiate projects that deliver cost savings while at the same time improving your organisations environmental footprint.

While it's right to focus on the data centre today, remember that it's only one part of the story.

Many of the data centre initiatives we recommend would be worth doing even if there were no concerns about the environment; server consolidation, improving the way you manage the heat produced by your machines and better resource management are all things that will save you money; And they are all projects you can start today.

Once you have done that, you can look beyond the data centre at your distributed computing environment.

Over the long term, you will be asked to help address the other components of the Bathwick Green matrix, so start thinking now.

Infrastructure 2.0

The final issue we'll look at is the notion of "Infrastructure 2.0". How will the increasing use of web-based software and Software as a Service (SaaS) impact your data centre?

Proponents of SaaS might argue that in a few years' time all you'll need in your data centre is a big, fast, network switch. Indeed, for smaller businesses this is a realisable goal today. Small businesses can buy email, hosting, CRM and accounting services as a network-delivered service now. While relatively few businesses rely completely on SaaS today, within five years a significant proportion will. We don't think that SaaS adoption will drive radical changes in infrastructure in the short term but over the medium and long term it will.

Organisations will increasingly use software that isn't running inside their data centre. So the need for reliable and secure networking will grow dramatically over the next five years.

Fundamentally the shift towards SaaS is about the distinction between utility services and differentiating services. Services like e-mail and back-up are utility services, in that they are only differentiating when they fail (when your e-mail goes down you are – unfortunately – 100% differentiated from your competitors who don't have e-mail problems).

Over the next five years, a growing number of IT services will be bought in as a network delivered service. You need to be ready to deal with the implications this will have for your network in terms of performance, security and reliability.

As organisations get smarter about what they should be focussing on, they will get better at buying in those services that aren't differentiating.

This paper was written independently by the Bathwick group for IBM.

**For more information please contact the author – Gary Barnett, Partner, The Bathwick Group
gary@bathwick.com**