

Er, what IS a data centre?

Introduction

Data centres are relatively new features on the UK's urban landscape. In fact, many people are unaware that they even exist. Often occupying boring-looking industrial buildings or nondescript offices, data centres take the form of rooms or halls packed full of computing equipment, patrolled by small numbers of geeky looking individuals who talk in a language only vaguely related to English. Nobody seems to understand why they are there or what they actually do, which in most cases appears to be shrouded in secrecy.

Yet data centres underpin an incredible range of activities across government, business and society and are now part of our critical national infrastructure whether we like it or not. So it seems worth finding out a bit more about them. This document explains in simple terms what data centres are, what they do and why they exist.

Section I	What is a data centre?
Section II	Why do we have data centres? What are they for?
Section III	Different types of data centre
Section IV	Why is the data centre business model so weird?
Section V	How to describe a data centre: Terminology
Section VI	How many data centres are there in the UK?
Section VII	Conclusion: Advice from Socrates
Appendix A	Fast facts: why data centres matter
Appendix B	Data Centre Jargon Buster



A purpose built data centre. Source NGD

Section I What is a data centre?

Here's the formal definition:

A data centre is a building (or self contained unit within a building) used to house computing equipment such as servers along with associated components such as telecommunications, network and storage systems.

A data centre is equipped with a guaranteed power supply and high bandwidth connectivity. Resilience is critical so redundancy (duplication) of networks, power and other infrastructure is common to ensure continuity. Other facilities include building management controls such as air conditioning to maintain the environmental conditions for the equipment within a specified envelope of temperature and humidity, and security systems to ensure that the facility and its data remain secure.

But what does this actually mean?

Data centres house computers so that they can

- a. do stuff like process, manage, store and transmit data and
- b. talk to each other or to digital equipment in other data centres or in offices, homes, vehicles, in satellites orbiting the Earth, on the moon or in fact anywhere you can think of.

Business processes, government services, telecommunications, transport infrastructures all depend on computers interacting in this way, exchanging digital information. Many of your everyday activities also rely on data centre processes, including obvious things like using your smartphone or the internet or sending email, and less obvious things like doing your shopping or catching a train. (See Data Centres: a Day In Your Life).

The kind of computers you find in data centres are known as servers (computers that are usually assigned to specific roles as opposed to personal computers which are more generic). Servers perform computing functions remotely from their operators and therefore don't have keyboards or screens or people hunched over them in the way that desktops would.



A server

Section I What is a data centre?


Servers run on electricity so data centres need to have a good power supply that will not get disconnected. This is called UPS or uninterruptible power supply. Back-up power supply is usually provided to make sure that even in the event of a power cut to the grid, the data centre is not left without electricity. Huge batteries provide instantaneous backup and there are usually generators on site with guaranteed fuel supply to provide for longer term outages. This is because so many of the critical functions of a modern digital economy rely on data centres that a power failure could have catastrophic consequences.

Servers also need to talk to other digital devices so data centres need to have excellent communications infrastructure in the form of high bandwidth connections (super-ultra-mega-fast-broadband), often from more than one telecommunications provider to ensure that if one option fails, connectivity is not lost. Essentially a digital economy is a networked economy and data centres can be seen as network nodes.

It might help to think of a data centre as a bit like a brain, which stores lots of information and runs both cognitive functions - the things we realise we are doing, and unconscious functions – the invisible things that we take for granted like breathing and keeping our hearts beating. If the oxygen supply to the brain is cut off then the consequences are rather significant. In the same way a data centre needs its own form of oxygen - a continuous, reliable power supply. And just like a brain has a nervous system to tell heart and lungs and arms and legs and eyeballs what to do, data centres need their own nervous system – telecommunications connectivity - to keep things like ATMs, street lights, telephones, hospital computers and air traffic control running smoothly (to name just a few). This is also the reason why data centres have to be highly resilient to external threats whether from terrorist attack or environmental cataclysm.



Generators: source Cap Gemini



Battery bank: source NGD

Simplistically, a data centre is a fancy shed filled with lots of computers talking to lots of other computers, most of which are elsewhere.

Section I What is a data centre?

What ISN'T a Data Centre?

An office is not a data centre, even if there are lots of computers in it. A call centre is not a data centre and neither is a disaster recovery site or a factory that manufactures computers or even a warehouse in which computer equipment is stored. Call centres and offices primarily accommodate people, a factory makes computers that aren't doing anything and things stored in warehouses are inert.

Your next door neighbour's garage, which he has filled with computers, isn't a data centre, even though he might like to think it is. Your next door neighbour's computers aren't adequately protected against changes in temperature or humidity or theft or flooding or power cuts or telecoms failures or his wife taking an axe to it all when resentment over her techno-widowhood overflows.

A car park or field or empty warehouse with planning permission and power supply is not a data centre either although, demand permitting, it could eventually become one and so could legitimately be described as "data centre capacity".

A data centre is there to manage, store, process, transact, manipulate or transmit digital data at scale, within a controlled, protected, resilient environment. If it isn't doing any of those things then it isn't a data centre.



Not a data centre

Section II Why do we have data centres – what are they for?

Data centres consolidate any number of separate IT functions within a single operating unit.

Twenty years ago there were no data centres – or at least none as we know them today. That's probably because there wasn't enough digital data to create a requirement for specialist facilities in which to house and process it, and the data that existed did not underpin enough critical government, business or social functions to make its integrity and security the key priority that it is today. The growing presence of data centres is the result of our increasing reliance on computing and on digital technology generally. More and more of our everyday processes, including government service delivery, business processes and individual activities rely on computing to function. The growth in data centres is also the result of changes in the way that we approach computing and how we choose to handle our computing resource.


Twenty years ago, big, complex computing jobs tended to be carried out by big, complex computers called mainframes that were housed in special facilities, usually in office buildings. Then as computing became more and more accessible, and individual processors became more and more powerful and capable, the model of "distributed computing" developed. Distributed computing just means that the computing functions are spread around according to where they are needed instead of being concentrated in one place.

More recently, as computing – and society – has become increasingly networked (see jargon buster on LANs and WANs) we need additional computers to manage these networks, collate and back up all the data from individual computers in one place, and manage applications like email, databases and internet presence. The computers that do these specialist functions are called servers and tend to be located away from people in server rooms or cupboards.



A traditional mainframe. Source not found

Section II Why do we have data centres – what are they for?



Growth in movement of digital data (Amsix)

During the last ten years two things have happened: Firstly the demand for digital data has grown explosively, (especially the demand for moving it around) and secondly, as computing has become more pervasive, we are increasingly reliant on it. If those computing processes fail for some reason or if data is destroyed or stolen the consequences could be catastrophic for business, government or individuals. Moreover there is the threat of punitive damages if this failure interrupts the provision of services to third parties.

These two factors, (rapid growth in the volume of data we need to deal with and our reliance on it) have made us reconsider the way we handle our computing functions. Locating servers in offices along with other corporate activities has the appearance of keeping all our eggs in one basket and could make a company very vulnerable to mishap. Unlike laptops or PCs, servers perform specialised functions so don't have to be close to the people who rely on them provided they are connected by a


telecoms network. So it makes sense to consolidate and protect these critical business functions in separate, specialised, secure facilities. Another benefit of consolidating computing in this way is that it is much more energy efficient than distributed computing.

The result of all this has been the rapid growth in data centres. Some companies have built their own, others use third-party providers. In fact a range of business models exists in this rapidly emerging market. These are explored in section III.

Section III Different types of data centre

Operational models for data centres: Colocation and Enterprise

The data centre market in the UK can be split roughly into two sub-sectors: colocation and enterprise. Colocation, or “colo” operators provide what is known colloquially as “position, power and ping” which basically means that they provide the infrastructure – security, constant electricity supply, broadband connectivity and an environment in which temperature and humidity is controlled to suit servers. They then sell or lease space within those specialised facilities to companies who install and manage their own IT equipment. The term colocation comes from the fact that these customers share or “co-locate” their IT operations in one purpose built facility.




Data hall in a colo site. Source: NGD.
The things that look like black wardrobes are cabinets containing racks of servers.

Colo operators are divided into wholesale and retail. Wholesale colo providers tend to lease whole sites or significant portions to their customers. These customers (in reality a wide variety of companies) in turn may sublet their share of the site or use it for their own corporate data needs or to service their own customers. Wholesale colo providers include DRT, Global Switch and Telehouse. Retail colo providers like Equinix and Telecity tend to lease space in varying quantities – from whole data halls down to a number of racks. Some colos offer customers a share in a single rack. Big colo data centres, like Telecity’s West London facility or Global Switch’s Docklands site, house the IT function for thousands of businesses.

Enterprise operators are those who use their data centre provision for their own purposes – i.e. for their own corporate IT functions (in the case of banks and supermarkets and government departments).

Enterprise operators also include those who provide IT services to third parties (HP, Fujitsu, IBM, BT, Atos, CapGemini, etc). Enterprise data centre operators may build their own data centre facilities or locate their data centres within colo facilities – Many companies do both. Some companies operate enterprise data centres but also sell some colo space to other companies, and this is a logical option for enterprise operators who find themselves with spare capacity.



Data centre module in an enterprise site. Source: CapGemini

Section III Different types of data centre

Why do enterprise operators use colo providers?

As mentioned above, enterprise operators may build their own data centres but for most it makes commercial sense to use a third party provider. Once an organisation's data requirements reach a certain size or become mission critical (where disruption of service has significant adverse consequences and generates liabilities), this data will need to be housed in an environment with guaranteed levels of security, continuity of power supply and connectivity.

Companies have a number of options here: They could build their own facility or they could take space from a wholesale or a colo operator and still manage their IT themselves or they could outsource the whole IT function to an IT services provider (who in turn may have their own data centre or have taken space within a colo). They could even buy a site and then contract a third party to manage their IT for them within it.


Companies that are not in the business of providing IT services are more likely to see data centre construction as outside their core business activity and use a third party and it is clear that many if not most enterprise operators take this view. Other reasons that encourage this approach include flexibility, minimising exposure to technological obsolescence, sharing risk and managing capital expenditure. This practice is so widespread that it would be hard to find a City institution or financial service provider that does not have a significant presence in one of the big London colos.

Large IT service providers may commission and build their own data centres from which they provide services to third parties. That is because IT services (and therefore data centres) **do** form part of their core business. HP for instance uses its large facilities in the North East to manage data functions for public sector clients (like government departments) and other customers. The number of these large facilities is limited and it is common for IT service companies to use colo providers for some of their offerings. So a large IT services company might only own two of their eight data centre facilities – the rest may comprise space leased from colo operators through a variety of models.

Section IV Why is the data centre business model so weird?

Expense and variable demand

Building a data centre from scratch is exceptionally expensive. Moreover, customer requirements vary enormously. Some customers just want a shell within which they can install all their own equipment. Some want a fully fitted space and others want a complete IT service. The expense and the variability in demand are two factors that have influenced the way that the data centre market has evolved.



A data centre?

In many cases (although not universally) data centres are not built speculatively. Land is acquired, the relevant planning consents are obtained, power provisioning is negotiated and then.... nothing happens. In some cases, nothing ever happens. This is because the developer is waiting for a suitable customer who will decide the way that the site is developed. So a company that says they have 10,000 square metres of data centre space available for customers may have a grassy space in a good location with the relevant consents, or a shell building with nothing inside it, or a partially fitted shell, or a complete data centre with lots of spare space. Now you can see why it is very hard to estimate how many data centres there are in the UK. A tally of “available” space will include a lot of grass, car parks or warehouses that may never become data centres.

Section V How to describe a data centre: terminology


Data centres are specialist, technology intensive facilities and the terminology used to describe them has evolved into a rather esoteric language (some might refer to it less kindly as jargon). This section explains the basic terminology used to describe data centres in terms of space, power and resilience, three very important characteristics used to compare facilities. In section VII there is also a jargon buster explaining some of the more common terms.

Terminology - Space

Data centres are described by size in terms of their gross space, their technical space and their core IT space (also known as white space or raised floor or net technical space). Gross space is the total footprint of the building in which the data centre is housed (sometimes even the total footprint of the whole site, car park and all, or the part of the site for which planning permission has been obtained), white space is the space reserved just for the IT equipment and technical space includes white space plus the supporting technical infrastructure (chillers, generators, network rooms and controls). Some operators view technical space and white space as the same thing and some use the term “data centre space” which could probably mean any of the above. So the definitions are not rigid. This unwelcome flexibility over terminology makes it hard to estimate the total area of data centre space in the UK.

White space is the most important measurement for a data centre because this is the space available to house IT equipment in the form of racks of servers. This is essentially the area that is available for lease or hire in colos so you can see why it is the most useful measure. In a data centre with 2,000 M² of gross space, only about 50% of that is likely to be white space. White space in turn is only about 70% of technical space.

As if this isn't confusing enough, data centre operators in the UK tend to mix units when describing data centres and use both metres and feet almost interchangeably. So a site might be described as 20,000 ft² of gross space with 10,000ft² of white space and with a power supply of 1.4 KW per square metre.



White space: Racks in a colo data centre. Source: TelecityGroup

Section V How to describe a data centre: terminology

Terminology: Power

Data centres are starting to describe their offerings not in terms of space but in terms of total power available to a site, which is known as power provisioning. Power availability is also frequently described in terms of the power available per square metre or per rack (expressed as KW per M² or rack or sometimes in Watts per square foot) or the total amount of power available to a site (expressed as MW or MVA, or for small sites, KW or KVA – see jargon buster for KVA/MVA).

A site with 10MW of power provisioning means that the operator has a guaranteed power supply and back-up provision of 10MW. The proportion of this power that they actually consume varies considerably from site to site depending on occupancy levels and the way it is operated. Don't forget that power in W – Watts - represents instantaneous power demand. Total power consumption is measured in Watt hours, Wh, or KWh, or MWh. One Wh represents a continuous demand of one Watt for an hour. 1000Wh is equal to 1KWh and so forth.

You might think that you could calculate the total power demand of a site by multiplying its power feed in MW by 8760 (24x365) to get an annual electricity consumption figure in MWh (MegaWatt hours). However, you would be wrong. The total power demand will never get anywhere near this figure because data centres are always heavily over-provisioned. When installing (provisioning) power, operators obtain the maximum power capability that they can. Firstly they have to provide power for cooling the servers as well as powering them and have to provide enough on the basis that the facility is full and that every server is working flat out and that it is the hottest day of the year and all the cooling equipment is flat out too. They will also be thinking about growth up to 20 years ahead. So you can safely assume that power consumption will always be much lower than provisioned power. For more information on this see [Data centres and Power: fact and fiction.](#)


Section V How to describe a data centre: terminology

Terminology: Resilience and security

Data centres are designed to help ensure business continuity. However, some data is more mission critical than other data so there is a grading system to indicate the level of resilience a data centre can offer using categories called tiers. This system has its drawbacks because it is based on American criteria for electricity supply which do not match those found in the UK. Tier 1 is the simplest and may just be a server room. The highest level is Tier 4, equipped to host mission-critical data and systems. There are only a few tier 4 data centres in the UK (to be honest there are not many activities that truly need Tier 4 resilience, despite what people may tell you).

Tiers 3 and 4 require top level security (usually with biometric entry controls) and uninterruptable power supplies – in which the back-up power supplies have their own back-up systems, and multiple broadband connectivity options involving a variety of routes and suppliers. Duplication of power supplies and connectivity to improve resilience is known as redundancy. Such data centres have very high levels of redundancy and there may also be a duplicate data centre facility at a separate geographic location to guarantee continuity in the case of an extreme weather event or other catastrophe. Operators also try to identify and minimise what they call SPOFs – single points of failure – to ensure that data centre function cannot be compromised by the failure of just one supporting system.

Anonymity and a general vagueness over location also contribute to resilience because companies may be keen to ensure that the precise coordinates of their mission critical data remain unknown. Data centre contracts between freeholders and operators or between operators and enterprise customers typically include strict non-disclosure clauses, especially if the customers are from the financial services sector.



Data centre security infrastructure: fences & cameras (NGD), PAC system (TelecityGroup), fingerprint biometric (CapGemini)

It depends....

This depends on what you are counting. If you are counting data centres as defined at the beginning of this document, then we are looking at about 300 recognisable facilities of significant size - say, over about 150M² of white space. If you count distributed computing and include server rooms and cupboards, then we are talking in the thousands, some say even tens of thousands. In terms of total power consumption, recognisable facilities probably use between 2TWh and 3TWh per year but again, this figure will increase if we include distributed computing.

Another way to quantify data centres is by the space they occupy, but again we have a problem. It is exceptionally difficult to differentiate “available space” which could mean anything from a field with planning permission to a fully fitted hall in an existing data centre, from “occupied space” which implies that the space is being actively used.

There are two other problem in quantifying data centre space in the UK: the Russian Doll problem and the “I’m Spartacus!” problem. These both result from the fact that a single data centre supports multiple layers of economic activity.

Russian Dolls

A single data centre can support a range of different business models all adding successive layers of service to the basic package – rather like Russian dolls. The little doll in the middle is the wholesale colo provider who may lease to either a retail colo operator or an enterprise operator. That retail operator will lease that space to a range of other operators some of whom may re-sell that space as “colo space” (these organisations set themselves up as colo providers but in fact they are re-selling space so are known as resellers). The lessees could alternatively use that space for their own servers and/or sell an IT service function to a number of other organisations. In this way a single data centre can provide IT functions for hundreds or even thousands of businesses. Some people are selling space, some people are selling services, some people are using the data centre to manage their own or other people’s applications which range from fleet logistics to web hosting.






"I'm Spartacus!"

The multiple activities that are supported by individual sites have also led to an "I'm Spartacus" tendency in the industry: each company with presence in a data centre will view that presence as "their data centre" even if in truth it is part of someone else's facility. Hence "we have five data centres" may genuinely mean five data centres but it could equally mean two data centres and a presence in three data centres owned by other people, or five halls in other people's data centres or a few racks scattered around five data centres or five racks in someone else's data centre or even five server rooms within their own office space. It can also mean five sites purchased speculatively for development or in various states of build.

The fact that a single data centre supports multiple service providers and multiple customers at the same time, all of whom regard that space as "their data centre" means that it is easy to over-estimate the size of the data centre sector in the UK.



Data centres are hard to characterise in traditional terms. They don't produce tonnes of output like manufacturing, they operate obscurely behind closed doors and high fences. They use technology that almost nobody understands that involves astonishingly complex processes. Their business models are weird and as a result, estimates of their size, number and power consumption vary widely. Ask a question about them and the answer you will get is "it depends" or "it's far more complicated than that..." followed by an incomprehensible explanation.

This paper does not aim to set out definitive answers on every aspect of data centre activity but it does establish some basic facts about data centres and explains why the sector in general is so poorly understood. At least then we know what we do know and what we don't know and what we understand and what we don't understand - and why. As Socrates said, ignorance is wisdom provided that you know what you don't know. So if, after reading these notes, you have learned that data centres are technically and economically complicated and that nobody really understands them, then this paper will have done its job.

The UK data centre sector is a spectacular success story: it has demonstrated rapid, sustained, and apparently recession-proof growth, without government assistance. In fact it is probably the only sector in the UK where the growth rate can compete with that of China. As the OECD stated in 2010¹ “Information and communication technologies (ICTs) are a key enabler of green growth in all sectors of the economy. They are a key part of government strategies for a sustainable economic recovery”.

Data centres also underpin and enable the digital economy so desired by politicians: Digital Britain, smart-grid and G-cloud cannot function without data centres. Data centres also support the very technologies that will reduce the carbon intensity of our economic activity – smart grid, smart transport, teleworking, e commerce and dematerialisation. Moreover, by consolidating computing activities into purpose built facilities, data centres are an efficient alternative to traditional ways of managing corporate and government IT functions.

Key facts

- Data centres underpin and enable the digital economy: they are its physical manifestation
- Data centres are the only geographical hook connecting the digital economy to a physical location
- A single data centre generates multiple levels of economic activity
- A single large data centre can house the IT function for thousands of businesses
- The UK currently dominates the European data centre market with around 60% of market share, spread between 250-300 sites with a combined power demand of 2-3TWh per year
- The increasing demand for digital data means that the UK sector is poised for further growth

¹ OECD, Greener and Smarter, ICTs, the Environment and Climate Change, September 2010

Challenges

- Data centres compete in a highly commoditised global market place
- Digital data is the most mobile commodity in existence
- Unilateral energy policy measures are putting the UK's competitive advantage at risk
- The risk takes the form of attrition as expansion and investment decisions cease to favour the UK and older facilities are decommissioned
- To maintain its dominance, investment in the UK should exceed £5bn but in 2012 was only £2.1bn
- Operators identify energy costs as a factor in their strategic decision making process

Data should be regarded as one of the key utilities of the 21st century, along with energy and water. In the same way that governments prioritise the efficiency and security of their energy supply and distribution, they must recognise in data centres a key technology, resource and skill set that should be retained within national and economic borders. All high tech manufacturing and knowledge economies are dependent upon ICT and on data centres in particular.

Appendix B Data centre jargon buster

The list below just deals with terms covered in this document.

More comprehensive glossaries of data centre terms are available from other sources. For instance the BCS and Green Grid have developed an excellent glossary of data centre terms: www.dcglossary.com/foswiki/Glossary/WebHome

Colo, colocation	A data centre or space in a data centre where power, connectivity, security and a controlled environment are provided and in which customers locate their own IT. Colo is short for colocation.
Collocation	Collocation just means someone has written colocation quite correctly and then Microsoft's autocorrect spellchecker has messed it up.
Enterprise	Refers to a data centre operated by an organisation to handle its own IT functions (in the case of banks and supermarkets and government departments). Enterprise data centres are also operated by companies who provide IT services to third parties
Reseller	Colo operators sell or lease space to other operators some of whom may re-sell that space as "colo space" (these organisations set themselves up as colo providers but in fact they are re-selling space so are known as resellers.) A reseller does not own the data centre it operates from – in simplistic terms it's a form of sub-letting.
Server	The kind of computer you find in data centres, without keyboard or screen, which tends to perform specialist functions automatically and remotely from its operator.

Appendix B Data centre jargon buster

Gross space	Gross space is the total footprint of the building in which the data centre is housed (or the part of the site for which planning permission has been obtained if the data centre is not yet built.) When people talk about site dimensions then they usually mean the size of the whole site including the car park, grass and decorative shrubs.
White space	White space is the space reserved just for the IT equipment and is the most important measurement for a data centre because this is the space available to house the racks of servers. This is essentially the area that is available for lease or hire in colos so you can see why it is a useful measure. In a data centre with 2,000 M ² of gross space, only about 50% of that is likely to be white space. White space in turn is only about 70% of technical space(qv).
Technical space	Technical space includes white space plus the supporting technical infrastructure (chillers, generators, network rooms and controls). Some operators view technical space and white space as the same thing. Others would say that net technical space is the same thing as white space.
Net technical space	The same as white space, also equivalent to core IT space.
Power provisioning	The total power available to a site, whether it is used or not. Power availability is also frequently described in terms of the power available per square metre or per rack (expressed as KW per M ² or rack or sometimes in Watts per square foot) or the total amount of power available to a site (expressed as MW or MVA, or for small sites, KW or KVA(qv)).
MVA: Mega Volt Amps KVA: Kilo Volt Amps	The short answer is the actual power as opposed to the apparent power. You can interpret MVA approximately as MW and KVA approximately as KW in this context. 30 MVA may end up being 29 MW but they are broadly comparable ² .

Appendix B Data centre jargon buster

Networks:

HAN: Home Area Network

LAN: Local Area Network

WAN: Wide Area Network

A Home Area Network (HAN) connects in-home digital devices, such as PCs, mobile phones, entertainment technology, thermostats, home security systems and smart appliances, into a common network. HANs can be wired or wireless.

A Local Area Network (LAN) supplies networking capability to a group of computers in close proximity to each other such as in an office building, a school, or a home. A LAN is useful for sharing resources like files, printers, games or other applications. A LAN in turn often connects to other LANs, and to the Internet or other WAN.

A Wide Area Network (WAN) is a network that covers a broad area (i.e., any telecommunications network that links across metropolitan, regional, or national boundaries) using private or public network transports. In essence, this mode of telecommunication allows a business to effectively carry out its daily function regardless of location. The Internet can be considered a WAN as well, and is used by businesses, governments, organisations, and individuals for almost any purpose imaginable.

2 The long answer is (take a deep breath): MVA is the product of multiplying the volts times the amps consumed and hence the electricity to be paid for. When the power is transferred 100% between the generator and the load (say a light bulb) the power consumed has the same value as the power generated: ie $MVA = MW$, we say there is a power factor of 1 or unity. This is a resistive circuit where the voltage and current are in phase. However when a load is powered by an AC generator from the electricity grid and the motor for say air conditioning is inductive, the power factor drops to say 0.9 or 0.8. This equates to the power being lost by a factor of 10 or 20%. This is because the voltage and the current in an inductive circuit are not in phase. This translates to a difference between the MW and the MVA figures. The power factor is the ratio of the real power versus the apparent power flowing in the circuit. The use of the MVA figure is used as it is the apparent power actually needed to run the air conditioner. The MVA figure is usually higher than the MW figure unless the power factor is unity. Inductors or capacitors are often used to correct an inductive load to bring it closer to a power factor of 1. Otherwise the company is paying for wasted power.

For more information please contact:

Emma Fryer, Associate Director, Climate Change Programmes
T 020 7331 2160 | E emma.fryer@intellectuk.org

James Harbidge, Programme Manager, Energy and Environment
T 020 7331 2173 | E james.harbidge@intellectuk.org

**Contributors and reviewers**

Text by Emma Fryer, Intellect

Reviewers

Jim Watts, DEFRA | Sylvie Feindt, DIGITALEUROPE | Alice Harrison, Intellect | Steve Hearnden, Intellect

Photos kindly supplied by:

Jonathan Tapp, CapGemini | Chris Williams, Next Generation Data | Felicia Asiedu, TelecityGroup

Lay readers:

Robert Burdall | Carole Fryer | Paul Vincent

Other publications in this series or discussed above include:

- Data Centres and Power: Fact or Fiction?
- Data Centres: A Day in YOUR Life
- So What Have Data Centres Ever Done For Us?

Generic publications relevant to this topic include:

- Data Centres: The Backbone of the UK Economy
- We Need to Talk About Jevons (final draft, publication imminent)
- Evaluating the carbon impact of ICT: The Answer to Life, the Universe and Everything
- High Tech: Low Carbon

Intellect is the go-to organisation for the UK tech industry. It is the trade association for the UK's tech sector, which represents developers and suppliers of digital technology and services. Intellect has over 860 member companies including major multinationals, mid-sized firms and small businesses. Collectively these companies directly employ more than half a million people in the UK and their products and services are used by every part of the UK and global economy.

© Intellect July 2013

The information in this paper is for use and dissemination on the condition that Intellect is referenced accordingly.