



FM, IT and Data Centres

Are Facilities and IT data centre managers implacable enemies, or is it just a need for different priorities and emphases on work that seem to get in the way?

January 2013

Often, Quocirca finds that an organisation has one team – facilities management (FM) – looking after the physical facility of the data centre, with another – information technology (IT) – looking after the servers, storage and network equipment, along with the software running within it. This can lead to problems where priorities clash, or where a lack of common language or views of a problem can stop things from happening. This report pulls together a series of articles written for SearchDataCenter throughout 2012.

Clive Longbottom
Quocirca Ltd
Tel : +44 118 9483360
Email: Clive.Longbottom@Quocirca.com

FM, IT and Data Centres

Are Facilities and IT data centre managers implacable enemies, or is it just a need for different priorities and emphases on work that seem to get in the way?

<Second précis>

Data centres can be run hotter than previously

A major means of saving energy in a data centre is to use less cooling – and new guidelines mean that a modern data centre can be run considerably warmer than previously. Combining this with other engineered approaches, such as hot and cold aisles, can provide large cost savings.

Originally published [here](#).

Facilities and IT must work more closely together

Facilities management (FM) and information technology (IT) teams are too often working in isolation – and this often leads to them working against each other. This must be addressed through combining teams and managing projects according to business priorities.

Originally published [here](#).

UPSs are critical components

The uninterruptable power supply (UPS) is no longer “just” a piece of equipment that is there as insurance. With increasing intelligence built in to the system, the procurement and use of a UPS has to be balanced with IT’s and the business’ needs – and FM and IT must work together on this.

Originally published [here](#).

DC power is no silver bullet

The hoary old chestnut of the use of DC power in a data centre keeps coming back. Although all IT componentry runs on DC, a DC-based data centre would be non-standard and expensive. Only those who will have massive, bespoke systems will find that a DC infrastructure makes economic sense to them.

Originally published [here](#).

There is increasing choice in how data centres can be implemented

“The” corporate data centre is not the right way to think any longer. Co-location and public cloud computing mean that a hybrid environment will be a more general means of providing an IT platform. For the corporate data centre, this does mean that changes will be required – and is a scale-out cloud or a modular approach best?

Originally published [here](#).

Physical security is just as important as technical security

Far too often, Quocirca has seen the focus being on creating massively secure technical platforms, with little focus on the physical security of facilities and equipment. FM and IT must work together in order to ensure that all security issues are dealt with to give higher levels of corporate security.

Originally published [here](#).

Conclusions

The data centre is changing. Although co-location, cloud and software as a service (SaaS) is moving some functions outside of the organisation’s control, the corporate data centre will remain a critical part of the overall IT platform for most organisations for the foreseeable future. The data centre cannot be seen as being in two parts – the facility and its contents – but must be designed, implemented, operated and maintained as a single, dynamic system. IT and FM have to work together to ensure that this is the case.



The changing face of ASHRAE data center environmental standards

When mainframes ruled IT, the received wisdom was to keep them as cold as possible. Water cooling was the norm, and cryo-cooling through the use of pumped refrigerants was Hollywood's preferred manner of showing supercomputers in use.

As the use of distributed computing spread, the interdependencies between the data centre facility and the computing equipment held within it became more complex. No longer was the main IT "engine" concentrated into one part of the facility – now, lots of smaller engines were spread around. These early tower systems still had to be kept cool, and many an IT manager has had servers collapse due to lack of cooling occurring when fans in such tower systems failed and systems management software failed to pick up the failure.

As the need for more compute power grew, the use of rack systems began to replace the use of towers. These standard-sized racks drove commoditisation of computer equipment into different multiples of height units (1U, 2U, 4U, etc.) within a 19" rack. Such concentration of equipment density made cooling even harder – radial fans gave way to axial fans, which can shift lower volumes of air.

The data centre facility itself became more important. Computer room air conditioning (CRAC) units became the norm, chilling and treating air to ensure that it could cool the equipment as required, without causing condensation through the moisture content of the air being too high, or the growth of dendrites which could cause electrical shorting through being too dry.

However, for many organisations, getting this right was a bit hit-and-miss, as no official guidelines were available as to what environmental conditions should be applied for cooling air within a data centre. To this end, the American Society of Heating, Refrigeration and Air-conditioning Engineers (ASHRAE) produces a document in 2004 laying out a set of best-practice guidelines as to what the environmental parameters should be for running a data centre.

In 2004, the design parameters of IT equipment were such that ASHRAE had to be quite prescriptive in its approach, and it also had to deal with predicted growth in equipment densities and thermal output from the equipment. ASHRAE could not depend on predictions of improvements in thermal and environmental envelopes of future equipment, however, which led to the advised parameters being well within the requirements of equipment launched even soon after the guidelines were produced.

In 2008, ASHRAE updated the guidelines to reflect that the pace of change in IT equipment had led to a different place than was expected – the increasing use of blades and of multi-core CPUs in multi-CPU chassis meant that equipment densities had massively increased, while the chip manufacturers had done much to improve both the thermal performance of their chips and the resiliency of them through the use of, for example, selective shut down of parts of the chips when not in use.

This second set of guidelines put the focus on maintaining high reliability of the equipment in a datacentre in the most energy efficient manner – a change from the 2004 guidelines which just focused on reliability. The increasing focus on energy usage within data centres means that measures such as power usage effectiveness (PUE) have become more important, and just maintaining reliability within a data centre without ensuring low energy usage is no longer valid.



To this end, ASHRAE has now expanded its data centre class definitions from 2 to 4 to provide a greater range of options to organisations where a better balance between reliability and energy effectiveness could be gained while still following best practices. As well as high-end enterprise class data centre guidelines, ASHRAE now covers server rooms and less mission critical environments. Provided that an organisation understands where its technical and business risks reside on its IT equipment, having four different sets of guidelines creates greater flexibility in the environmental choices for different parts of a data centre.

The 2008 document gave general guidelines of a dry bulb data centre temperature of between 15 and 32°C (recommended 18-27°C), with an allowable relative humidity range of 20-80% and a maximum dew point of 17°C at a maximum elevation of 3050m for a class 1 data centre. The allowable ranges for a class 2 data centre were marginally greater, but the recommended levels were pretty much the same. These guidelines extended the upper limits of recommended temperature limits by 2°C, and the upper recommended limit for relative humidity by 5%. This could have had a major impact on data centre cooling costs – except the majority of data centre owners still preferred to use the “carbon life-form guidelines” – i.e. keeping the data centre at a temperature more suited for employees, or around 21°C.

The 2011 guidelines still retain the same overall recommended temperature ranges, relative humidity and elevation across the 4 types of data centre. However, the allowable operating ranges now cover from 15-32°C for enterprise-class servers (an A1 data centre), through to 5-45°C for volume servers, storage, PCs and workstations (an A4 data centre). Therefore, by “zoning” equipment correctly, a data centre can be run under mixed environmental conditions, and energy usage around cooling can be minimised, and often almost eradicated through the use of free air cooling or other lowest-cost approaches.

It is good to see that ASHRAE is keeping the guidelines dynamic and reflecting a balanced view between the needs for reliability and energy efficiencies. When combined with facilities best practices in building and use of low-cost cooling approaches, along with equipment best practices in the use of hot aisle/cold aisle and complex fluid dynamics (CFD) hot spot identification and eradication, an ASHRAE-class data centre should enable IT and facilities management to ensure that the technical platform along with the physical data centre meets the needs of the organisation. However, data centre managers will have to realise that the data centre is built for silicon-based systems, not carbon-based ones – and that cooling to 21°C is just throwing away good money.

Facilities management team to the IT department: Let's work together

The data centre is a building – and facilities management tend to have this under their control. However, they also have every other building within the organisation under their control as well. With a diverse portfolio from warehouses to offices, kitchens to restrooms, car parks to storage systems, the data centre can just be another item to them.

For the IT department, the data centre is the centre of its universe. The main components that support the mission critical aspects of the business are the servers, storage and network systems that are housed within the data centre – and if anything goes wrong, then it is the IT department’s neck on the line – even if the problem was due to a facilities issue.



What happens when there is a mismatch of priorities? A blocked executive restroom could be upsetting a lot of people, and the facilities management team may prefer to put a few people on this to get rid of the complaints of senior management, whereas a request to just “check up on the cooling in the datacentre, please – we’ve noticed a small rise in temperature” may not be seen as being quite as important.

That the executives could use a different restroom and that the small rise in temperature denoted a failure of a cooling system that is just about to lead to the enforced shutdown of all the systems in the data centre just didn’t really come into the equation – and this is why facilities and IT have to work more closely together.

Does this mean that facilities and IT must start having long interminable meetings to discuss matters that are seen as being critical to one side and not quite so critical to the other? No – what it does mean is that each side must be able to see what is the root cause of any problem rapidly as and when required – and this means that as much as possible needs to be automated.

The key to this is using monitoring, measurement and management systems that are common, sharing the same data sources. This way, the building information system (BIS) used by facilities is plugging into the exact same data used by IT’s systems management tools. When IT sees a small rise in temperature on just one piece of IT equipment, it’s likely to be just a fault in that piece of equipment – but if they see a rise in temperature across the board, it is far more likely to be a facilities issue. Likewise, hot spots that could lead to a risk of fire are better served from the use of facilities type systems such as smoke and infrared detection systems.

A relatively new approach to bringing the tools and systems closer together is now being used by some organisations. Data centre infrastructure management (DCIM) started off as a specific set of tools for facilities to look after the data centre as a special case. DCIM tools tended to enable facilities to look at how the data centre was laid out, how power distribution and cooling was configured and so on. As its usage grew, other functions were introduced, such as computational fluid dynamics (CFD), structured cabling, and this then led to areas such as granular power management with some DCIM tools coming with comprehensive equipment databases covering real-world data on how much power a specific server, storage system or network device would draw.

As the DCIM systems became more comprehensive, they began to be seen by the IT department – and much of what the tools did was seen to be of use to them. Being able to see exactly how much power a rack or a row would draw meant that IT could look to see if the data centre was capable of providing this. Using CFD allowed them to ensure that their usage of racks and rows did not lead directly to poor cooling flows.

Where space limitations force IT to make decisions that could lead to a need for more power or better cooling, using DCIM tools means that facilities can be included in the decision – is it possible to bring in more power just here, or can the amount of cooling be improved through better ducting to this area?

Further, facilities’ BIS systems may well be a good feed into DCIM as time goes on. The move to “intelligent” buildings has been slow, but moves to control carbon emissions will drive the use of IT across the organisation to more optimally manage areas such as heating and cooling, along with the security of buildings and the management of control systems on the shop floor. The only way that such a complex mix of technologies can be effective is through everyone involved having a true view of everything that is happening. At this stage, not only will facilities and IT have to ensure that everything is working closely together – but they will also have to ensure that dashboards and reports are suitable for the business as well.

At the most basic level, it is time for IT and facilities management to look at how DCIM can help both their teams. At a more advanced level, main boards should be looking at how DCIM can help in driving energy and carbon savings across their whole organisation – while providing higher availability and optimal running of their critical IT systems.



Does UPS maintenance go to IT or the facilities management team?

An uninterruptible power supply (UPS) has one simple job to carry out – to seamlessly and immediately take over the task of providing power to a data centre when the main supply from the grid is interrupted. As such, it has been seen as being part of the fabric of the data centre facility, and has fallen under the auspices of facilities management (FM). After all, the procurement and management of the grid-based energy supply is FM's – so why not the UPS?

The issue is that the world within the data centre is changing. Energy costs are highly variable, but trending inevitably upward. Equipment densities are increasing, with many racks now drawing in excess of 15kW.

In a paper published in 2010, Professor Jonathan Koomey calculated that power usage across data centres worldwide had increased by 56% between 2005 and 2010 – less than the originally expected 100%, but still taking usage up to around 26GW of energy being needed to power data centre facilities – or the output of 26 average-sized power stations.

At an average power usage effectiveness (PUE) of 1.8, this means that the energy used by non-IT equipment is around 14.5GW. The majority of that is in cooling down data centres, but a large part is also in running UPS systems.

A typical UPS system consists of a mechanism for providing stored energy – generally through the use of rechargeable batteries. Such systems will be able to provide continuous power for only a matter of minutes or, at best, hours, and so need to be backed up by alternative sources of power generation – which tend to be petrol or diesel-powered generators.

Although a UPS is nominally “off” for most of the time, the energy impact of a poorly implemented or managed system can be massive. Older systems would take a continuous trickle drain from the main grid supply in order to ensure that the batteries were kept fully charged. Over-charging batteries can lower their effective life, and so many UPS had complicated systems to manage the power being drawn – but these could also be poorly effective, so resulting in lower energy efficiencies. In order to optimise the batteries' life, they would need to be deep-drained on a regular basis, requiring either the data centre to be run from the UPS until the batteries were sufficiently drained, or shunting the output to a false load – in each case, introducing a weak link in that if the grid power failed, the UPS would have less capability to kick in for long enough for the generators to take over.

Also, a petrol- or diesel-powered generator will need the fuel either draining and replacing on a regular basis (to prevent the fuel going “stale”), or the system running on a regular basis to use the fuel up. With smaller diesel generators being far less energy effective than centralised main power generators, such a fuel cycling can be both costly at a financial level, and hit an organisation's green credentials hard. However, running the data centre via the UPS and generator does, at least, provide a test for how effective the power backup procedure is.

Another issue is that as energy densities build up in the data centre, there becomes a need to also introduce more UPS capability to match the growth in data centre need. This has driven a move from the monolithic UPSs of old to a more modular approach, with a corresponding growth in use of in-rack and in-row UPSs such that incremental needs can be met as new racks and rows are introduced.



However, at the generator level, this is not so easy. Sourcing incremental capability and ensuring that the generator's output is matched with other generators' output is not always easy – and most data centres choose to go for a forklift upgrade – replacing the generators completely – as the need arises.

But, tying the data centre equipment energy usage, and the priorities around the workloads on the equipment, can lead to a more intelligent and useful UPS strategy.

Although an organisation is increasingly dependent on IT, it is not equally dependent on all of the components of its IT. Whereas a retail organisation losing power to its ecommerce web site could lead to the failure of the company, losing power to servers that support its payroll is not so much of an issue. Using the intelligence built-in to modern equipment can lead to cpu cores being powered down when not in use, disk drives being spun down and low-energy states being entered where necessary.

Forming the strategy around which workloads to fully support and how low-energy states can be used to maintain different levels of IT capability is not an FM task. IT has to work with the business in order to prioritise workloads and what levels of capabilities have to be maintained over a period of time. The use of integrated in-rack/in-row UPS systems provides the granularity required for IT to be able to create an optimised power usage strategy. However, the technicalities behind modern UPS systems – such as power factor control (PFC) and total harmonic distortion of the input current waveform (THDi) are best suited for the FM team, which needs to understand how the UPS system will fit in with the broader energy strategy and needs of the overall organisation.

As is to be expected, the answer to the original question – is a UPS an FM or IT matter? – is that it is both. FM and IT have to work very closely together to ensure that any system put in place is fit for purpose. It must support the organisation in maintaining critical IT workloads, while also minimising overheads and energy wastage.

Is direct current power the silver bullet for data center efficiency?

The vast majority of electrical components (as in chips, resistors, capacitors, etc.) in a data centre run off low-voltage direct current power. Yet, we insist on bringing in high-voltage three- and single-phase alternating current energy, which has to run through multiple different step down transformers until the desired voltage is reached. Each transformation involves a loss of energy – even the most modern transformers will generally be no better than 98% efficient.

Consider the power distribution path to your data centre and then within it. Firstly, the power generated at the power station will need to be boosted up to a high voltage suitable for transmission. This will get to a substation, where it will be brought down to a distribution voltage. There may be 2 such steps involved, after which it will reach the data centre itself, at a 450V three phase or a 110/220/240V single phase supply. From here, it needs to be stepped down to the main 12V, 5V, 3.3V and 1.5V DC voltages that are used by the IT components. This shows a straight path of around 4 transformations – or an overall efficiency of around 92% - 8% of the generated energy has been lost purely through transformers – and this is only where the transformers are 98% efficient.

In reality, it is far worse than this. The power distribution within a data centre will tend to be single phase AC. Each item of IT equipment will have a main transformer inside it with multi-voltage outputs, with many more on-board transformers to provide the voltages required by specific components. However, larger transformers tend to be more



efficient than smaller ones – so these mini-transformers may only be 95% efficient. If this is the case and 6 steps of transformation are involved, then over a quarter of the generated energy is lost in transformation.

Therefore, it is surely good practice to use direct current across as many areas as possible, avoiding the need for all these transformers in the equipment within the data centre. However, there are a number of issues that need to be taken into account to see whether this is really possible.

Firstly, at a power transmission level, DC is highly inefficient. Through the laws of physics, to get a set amount of power (in Watts) down a line is voltage (in Volts) x current (in Amps). High transmission voltages in the hundreds to thousands of kiloVolts are used to lower the current. This is important, as resistive losses are linked to the current. DC or AC can be used at high voltages for power transmission, but AC has been the main choice to date – and so trying to lose the transformation losses for the transmission stage is unlikely to be workable. It's pretty much the same at the distribution level – changing the existing infrastructure to move from AC to DC is not going to be easy, and there would still be a need to use high voltage DC, so still requiring transformation along the way.

Those who advocate the direct current data centre point to the telecommunications industry, whose facilities have run using DC for a long time. There are reasons behind this, however – a full industry was built up around the provision of DC equipment, and there was less need for multiple DC voltages within an old telecoms facility. Therefore, power distribution within such a facility could be carried out using transformers built into the fabric of the facility and then large copper buss bars to enable high current DC to be distributed. Also, in the 1980s and 1990s, although variations in oil prices were causing some uncertainty in energy pricing, the main generation capability was still through relatively cheap and available coal, meaning that energy efficiency was not such a focus as it is today.

The main issue for a DC focused data centre is that the equipment is still not widely available. Although the components are DC, the vendors have to focus on the mass market, and so build equipment that uses AC as its main input. Some vendors do DC versions of their equipment, but the lack of uptake means that these are premium priced. With the additional costs of kitting out a facility with DC power transformers, DC power distribution and management systems, using servers, storage and network systems that perform to all intents and purposes the same as the AC variants but cost more just does not make sense. Better to follow the crowd and make use of the lower costs of standardised, AC-based systems.

There are two big hopes for the DC data centre, however. One is the modularisation of the data centre. As systems such as Cisco's UCS, Dell's vStart, IBM's PureSystems and others come through, these pre-configured "blocks" can be wired internally in any way the vendor wishes. It makes sense for the vendor to cut out unneeded componentry, and multiple transformation stages can be cut out during the design and build phase. In the same way that cooling systems are moving from the facility to the module with in-rack cooling systems, it becomes far more likely that power management will move from the facility to the module as well, with in-facility power distribution being based just around the provision of single-phase AC cabling to the point of need.

The second is the growing prevalence of cloud computing. For a cloud provider with tens to hundreds of thousands of servers, massive storage infrastructures built around scale-out storage units and a complex network structure, buying in DC from the start could create a viable payback period.

Therefore, Quocirca's advice when it comes to the DC data centre is: don't worry about it. A strategic decision to move to a DC infrastructure is likely to end up expensive and forcing the organisation to be dependent on certain types of hardware. However, the vendors will drive a better energy efficient direction through optimising the DC usage within their own systems.

As an organisation, you will get more DC in the data centre – it's just that it will happen of its own accord.



Flexible computing in business: Move to cloud computing or go modular?

Depending on who you listen to, the impact of cloud will lie somewhere on a spectrum of “very little – I wouldn’t touch it with a bargepole” to “complete game changer – this signals the end of the IT department”. Just where along this spectrum does reality really lie?

Unfortunately, it’s not that easy to provide a definitive answer, as the correct place for your organisation along this spectrum is dependent on so many variables, such as your own organisation’s risk profile, the preponderance of in-house applications and the age of an existing data centre facility and its contents.

When looking to create the optimum private data centre facility for IT, it is best to start from a viewpoint of flexibility being everything. It makes no difference if cloud is going to be the biggest thing to hit your organisation or not, any lack of flexibility in the data centre could be catastrophic to the business.

If we start from a position of *laissez faire* and everything continuing “as is” (i.e. everything in the data centre, and little change to the IT architecture), the data centre still has to be able to grow as the business grows. If the current facility becomes a constraint, then it will have to be replaced – and this is not the best option under current financial conditions. Therefore, something else has to be looked at to provide the flexibility that IT and the business demand.

In comes virtualisation, as a means of lowering the amount of equipment needed to manage current workloads – often by 50% or more. This sharp drop in the amount of IT equipment under management may sound great, but if this is all housed in the middle of the same sized facility without any changes to uninterruptable power supplies (UPS), backup generation and cooling, the data centre will not be optimised – and power utilisation effectiveness (PUE) values will climb through the roof.

The counter to this is that virtualisation is a once-only “fix”. Once it is in place, growth of IT equipment and the space it needs is then likely to occur again – unless new workloads are pushed out of the facility into the public cloud.

So just what does the use of cloud mean to how the existing facility is architected? Cloud computing is still in its early years, and although the theory is strong, in many cases, the practice leaves a little to be desired. Therefore, architecting a data centre such that it shrinks as workloads are pushed out is one thing, but if the workloads are to be brought back in due to a cloud provider not meeting requirements or going out of business, how easy will it be to grow the private facility rapidly to embrace this? Even where the cloud provider meets its responsibilities, what else needs to be done within the facility in order to ensure that the end user experience is good enough for continued use?

The first thing is to move to a more modular off-the-shelf IT equipment model, using the likes of Cisco UCS, VCE V-Blocks, IBM PureFlex or Dell vStart systems, rather than build-your-own racks. Although build-your-own may sound like it provides greater flexibility, actual speed of response to business needs is often compromised, whereas the pre-configured modules in modern systems can be put in place and provisioned rapidly and more resources can be added as required in a very effective manner.

Greater modularisation also enables the use of a more structured approach to other areas of the data centre. Using hot and cold aisles in the private data centre allows more targeted cooling, and when used along with higher temperatures and variable speed CRAC units (or even free air cooling) as well as resized and more granular UPS and



generation systems from the likes of Eaton, Emerson or Schneider allows greater control over how the facility and the IT equipment work together.

To optimise the usage of space within the facility, false walls can make it that spare capacity can be walled off for use as office or other business space – but it should be remembered to ensure that these walls do fit from sub-floor to raised ceiling to stop any leakage of cooling air through gaps above the dropped ceiling or beneath the raised floor. However, the capture of hot air from cooling systems can be easily used within colder climates for space heating in these walled-off areas, or can be improved using heat pumps for example for water heating in hotter climates.

In order to ensure that the end user experience is optimised for a hybrid cloud environment, it is recommended that wide area network (WAN) acceleration or optimisation systems are investigated, so that the latency across a much more distributed IT platform is minimised. Here, vendors such as Silver Peak, Riverbed and Cisco offer hardware and software services that can ensure that performance of applications is maintained across long distances.

If cloud computing is being introduced within the private data centre, then it will also be important to ensure that applications are architected correctly. Any network traffic should be kept as much within the facility as possible. This can be done through the use of virtualised desktops, with the business and presentation logic all being essentially co-located, and only the visual aspects of the interaction being presented to the user's device. Not only does this ensure the best levels of performance, but also enhances security through keeping all storage centralised.

Embracing cloud computing means that IT and the facilities group must work more closely together. Creating a cloud architecture without addressing the facility will result in PUE values that could negate the capability for an organisation to claim that it is aiming for sustainability in its computing. For a facility to be “built” for cloud without understanding the dynamics of the IT involved and the likely strategy for outsourcing and insourcing workloads will result in a data centre that does not have the flexibility to truly support the business.

On the spectrum of adoption, cloud computing is likely to cause a continuous downshift in the focus of computing within a private data centre – but only over a period of many years. Bring IT and facilities together will ensure that the spectrum can be travelled along, at the right speed and capabilities to support the business.

IT security issues: The perception and deception of security

For many organisations, IT security is a top-of-mind subject. Organisations look at multi-factor authentication in order to try and ensure that the end user is who they say they are; tokenisation to try and stop “man-in-the-middle” hijacking of sessions; encryption of data on the move and at rest plus the use of VPNs for all sessions outside of the firewall as well as anti-malware systems such as anti-virus and distributed denial or service attack throttling. All are relatively common approaches to securing the organisation's data and its availability to the right users. Many organisations have all of this in place, so are they secure?

Unfortunately, a pure IT security approach can lead to a problem Quocirca calls a “perception of security”. It is all well and good creating a Fort Knox approach to IT security, but if physical security is more like a Gruyere cheese, the business is not secure and intellectual property can still leak out all too easily. As an example, look at the politicians who carry a sheaf of papers that can be easily photographed by any press photographer in the vicinity.



An organisation needs to create a risk profile that it can form a full security policy around. In many cases, security issues within an organisation are actually caused by accidental information leakage by employees or others working with the organisation, such as consultants and contractors. At a technical level, this can be dealt with through the use of tools such as data leak prevention (DLP) and digital rights management – which can also help to head off some of the more malicious attempts to redirect information.

However, this still leaves the less technical areas of information security. Facilities need to work with the business in order to secure other possible sources of weak security that may be more under their control, such as printers, where “pull printing” (using PIN codes or tokens to release a print job at a specific printer) can ensure that confidential papers are not left lying around or are taken by the wrong person. Fax machines should be replaced with multi-function devices that feed into the standard IT environment using scan to fax technology so that DLP systems can be brought to bear on these as well.

Even at a telephony level, it should be possible to put in systems that record a proportion of calls – provided that both the employee and the caller are aware of the fact of the possibility of being recorded. Call recording cannot be used for dynamically cutting off a call, but can at least act as a deterrent for malicious security breaches, and can also be used for training purposes to show how a simple call can put at risk intellectual property through a slip of the tongue. All of these require facilities and IT to work closely together, as technology is brought in to help ensure the overall corporate security policy is more easily enacted.

Organisations need to write into their contracts of employment what the physical security policies are, and what wilful disregard of these policies means to the employee. For example, an organisation should retain the right of searching an employee and their bags at any time, and opening any physical mail that is sent from within the organisation to an external address. All items containing corporate data and information must be returned on the employee leaving the organisation or the data held on a device must be provably destroyed if the device is under the ownership of the employee through, for example, a bring your own device (BYOD) system. Even at the extremes of securing intellectual property where (as yet) technology cannot be brought to bear, employees and partners need to understand that using knowledge that is in their heads against the business will result in possible sanctions – right through to criminal proceedings where necessary.

At the data centre level, having great technical security still leaves the problems of physical security. Facilities can provide better data centre security through the use of, for example, anti-ram-raid bollards or barriers and high-quality closed circuit television (CCTV) to put off those wanting to break in. This should be backed up with no windows in the facility – this makes breaking in harder, but also stops anyone just looking into the facility and makes it harder to use vibration detection to pick up voice or RF detection to pick up electrical signals. Electricity and internet cables should be armoured as well to prevent malicious physical denial of service attacks.

Multi-level security systems can then be applied for those who should have access to the facility: multi-factor security on entry doors, maybe using biometrics and/or one-time passcodes, and tracking capabilities within the facility based on, for example, near field communication or smart card systems. Tracking of people via CCTV again helps in ensuring that “tail gating”, where two people go through a door with a single security pass, is avoided. Physical cages within the data centre operated by electronic locks ensure that only those who are authorised to access specific machines can do so – and contract engineers brought in to carry out work can be given time-limited access to the physical systems.

Total corporate security is not something that can be done through IT alone. This requires a mix of business-driven policy supported by physical and technical security operated in a seamless manner. To attain this, IT and facilities



have to work more closely together – or find themselves as the scapegoats when security is breached and the business demands to know why.



REPORT NOTE:

This report has been written independently by Quocirca Ltd to provide an overview of the issues facing organisations seeking to maximise the effectiveness of today's dynamic workforce.

The report draws on Quocirca's extensive knowledge of the technology and business arenas, and provides advice on the approach that organisations should take to create a more effective and efficient environment for future growth.

About Quocirca

Quocirca is a primary research and analysis company specialising in the business impact of information technology and communications (ITC). With world-wide, native language reach, Quocirca provides in-depth insights into the views of buyers and influencers in large, mid-sized and small organisations. Its analyst team is made up of real-world practitioners with first-hand experience of ITC delivery who continuously research and track the industry and its real usage in the markets.

Through researching perceptions, Quocirca uncovers the real hurdles to technology adoption – the personal and political aspects of an organisation's environment and the pressures of the need for demonstrable business value in any implementation. This capability to uncover and report back on the end-user perceptions in the market enables Quocirca to provide advice on the realities of technology adoption, not the promises.

Quocirca research is always pragmatic, business orientated and conducted in the context of the bigger picture. ITC has the ability to transform businesses and the processes that drive them, but often fails to do so. Quocirca's mission is to help organisations improve their success rate in process enablement through better levels of understanding and the adoption of the correct technologies at the correct time.

Quocirca has a pro-active primary research programme, regularly surveying users, purchasers and resellers of ITC products and services on emerging, evolving and maturing technologies. Over time, Quocirca has built a picture of long term investment trends, providing invaluable information for the whole of the ITC community.

Quocirca works with global and local providers of ITC products and services to help them deliver on the promise that ITC holds for business. Quocirca's clients include Oracle, IBM, CA, O2, T-Mobile, HP, Xerox, Ricoh and Symantec, along with other large and medium sized vendors, service providers and more specialist firms.

Details of Quocirca's work and the services it offers can be found at <http://www.quocirca.com>

Disclaimer:

This report has been written independently by Quocirca Ltd. During the preparation of this report, Quocirca may have used a number of sources for the information and views provided. Although Quocirca has attempted wherever possible to validate the information received from each vendor, Quocirca cannot be held responsible for any errors in information received in this manner.

Although Quocirca has taken what steps it can to ensure that the information provided in this report is true and reflects real market conditions, Quocirca cannot take any responsibility for the ultimate reliability of the details presented. Therefore, Quocirca expressly disclaims all warranties and claims as to the validity of the data presented here, including any and all consequential losses incurred by any organisation or individual taking any action based on such data and advice.

All brand and product names are recognised and acknowledged as trademarks or service marks of their respective holders.