

## It's all change at the DCA

By Steve Hone CEO and Cofounder, DCA Trade Association



IT'S ALL CHANGE at the DCA for the start of 2017. There are now two new staff at the DCA to help and support members.

Kieran Howse joined the team at the start of the New Year as Member Engagement Executive replacing Kelly Edmond; he is already up to speed and here to support you moving forward.


I am also pleased to announce that Amanda McFarlane has joined the team as Marketing Executive. Amanda has extensive marketing experience in the IT sector. This new role has been created to ensure information continues to be effectively disseminated to our members target audiences; be that government, end users, general public or supply chain. Amanda has already taken over responsibility for the delivery of the new DCA website and members portal which is due to come on stream very soon.

Many of you will have had the pleasure of both working and spending time with Kelly Edmond over the past three years and I

know you will join me in wishing her all the best for the future and extending a warm welcome to our new staff.

This month's theme is 'Industry Trends and Innovations'. I would like to thank all those who have taken time to contribute to this month's edition allowing us to start the year with a good selection of articles.

During a recent visit to Amsterdam to attend the EU Annual COC meeting and the 5th EU Commission Funded EURECA event. I had the pleasure of getting a lift in the New Tesla X, not only does it run on pure battery power and do 0-60 in 2.6 seconds (quicker than a Ferrari) it literally drove itself, which when you think about all the innovative technology which must go into this car, the computer power needed to ensure it remains both reliable and safe - it's quite mind blowing! If this is a barometer of the sort of trends and innovations we can expect to see within the data centre sector, then we certainly have exciting times ahead. The DCA plans to dedicate a section of the new website to showcase our members ground breaking



ideas as they emerge. Please contact Kieran or myself if you wish to spread the word about something new we'd be very keen to hear your news.

The theme for the March edition is 'Service Availability and Resilience', the deadline date for copy is the 23rd February. This will be closely followed by the April edition which this year offers members the opportunity to submit customer case studies, the copy deadline is the 21st March.

Looking forward to a productive year, thank you again for all your article contributions and support.



## Data centres as a service industry



THE BENEFIT of hindsight is that it provides clarity and perspective. It is now clear that - mergers, acquisitions and ground breaking technologies aside - there was an obvious shift in thinking last year towards the data centre now being regarded as a utility. Going forward, that most likely means that there is also going to be a much greater emphasis than ever before on service excellence in the data centre in 2017.

The key driver behind this is that data centres are becoming increasingly core to business operation and success, and are no longer considered to be 'simply running behind the scenes' but ensuring backup and storage is managed effectively. Information continues to be the lifeblood of global businesses, the mission-critical data centre is now deemed to be an asset that needs to be protected at all times. Data centres are playing an increasingly important role in providing the essential infrastructure required to support 'always-on', content-rich, connectivity-hungry digital international businesses, we must perhaps now expect that the way in which they are procured will also change.

What is shifting first and foremost is that the major purchasers of data centre space are now more likely to be hyper-scale companies such as Amazon, Google and Alibaba, than large corporates. The move to the Infrastructure-as-a-Service (IaaS) and Software-as-a-Service

(SaaS) model, or any other such cloud services for that matter, will only continue to grow. The hyper-scale businesses know what they want and exactly when they need it. Their expertise in providing a hosted service to their customers means that they are 100 percent focused on efficiency, pricing, flexibility and reliability of service from the data centre operators that they ultimately select as business partners. They set the business agenda and the data operator is expected to meet it.

As a result, the requirement for consultancy and 'old fashioned outsourcing' from data centre operators is being replaced by a smart new approach, driven by an overall preference for finding a different way of working that focuses on taking a far more collaborative approach.

Data centre operators will without doubt continue to be selected based on their ability to build a facility that is sophisticated but added to future demands and expectations will be the underlying ability to form a trustworthy partnership with their clients.

The focus will become about showing detailed levels of understanding about the specifications provided by the client's highly knowledgeable and experienced engineering and IT team, whilst delivering an agile yet cost effective implementation and maintenance model in the longer term.

As a service industry, data centre operators will therefore need to expand their levels of knowledge to identify lower efficiencies and continued cost savings. It will also require a responsive 'can do' attitude, complemented by advanced problem-solving skills and a heightened awareness of what guaranteeing a 24x7 operation actually means.

The impact of globalisation also suggests that data centre operators should be prepared to be asked to support client expansion into new territories, practically on demand, by harnessing their own experience of what it takes to enter new markets successfully, and then build quickly.

Extensive advanced planning, combined with the ability to deliver the same infrastructure in more than one country - against increasingly short deadlines - in order to support ambitious go to market strategies will become critical requirements in the RFP.

Looking ahead, we fully expect data centre operators to have to work much harder to meet the needs of fewer but larger customers and our flexible business model means we are ready. For those intending



to deliver a new kind of data centre service in 2017, agility and adaptability will be what keeps you ahead of the game. Don't underestimate the importance of this.

[www.zeniumdatacenters.com](http://www.zeniumdatacenters.com)

## Weathering the perfect storm.

### Why innovation should not just be about the 'tech'

By Dr Theresa Simpkin, Head of Department, Leadership and Management, Course Leader Masters in Data Centre Leadership and Management, Anglia Ruskin University



WHEN TIME MAGAZINE named the personal computer the "Machine of the Year" in 1982 few would have accurately predicted the rampant

technological advances that stitch together the fabric of our economies and underpin our personal interactions with the world.

The introduction of the internet and all that has come with it, the explosion of big data, innovations in artificial intelligence and machine learning have brought with them profound change in an extraordinarily snapshot of time.

The tools available to us to be better at what we do (be it related to an IoT connected golf club that improves one's swing or making better medical diagnoses based on advice from AI applications) have evolved exponentially. However, the way in which we work is still largely based on post war habit and leadership capabilities that have changed little in the last half century. It is clear that despite sensational advances in technology and the applications to which they have been put, we are yet to apply

similar attitudes to technical innovation and change to the human and organisational structures that form the spines of our businesses. It is ironic that the Data Centre sector is as much a victim of outmoded business thinking as many other traditional sectors.

There is a raft of evidence to suggest that the sector is facing some big ticket issues; ones so large that individually they put business objectives at risk but together they form a perfect storm that dampens growth and may limit carriage of strategic objectives such as profitability and growth in market share. One of the most pressing is the discussion regarding the capabilities inherent in the human aspects of this 'second machine age'; the rampant advancement of technology bringing with it seismic shifts in the way we work, the skills we develop and the capabilities we need to take organisations and economies into a space where there is optimal capacity for efficiency, profitability and public good. For example, there is a raft of research from industry associations, vendors and business analysts that identifies skills shortages as a pressing and immediate risk to business and the sector as a whole.



Shortages in occupations such as engineering, IT and facilities management are leading to a rise in salaries as the 'unicorns' (those with a valuable suite of capabilities) are poached from one organisation to another driving up wages and generating a domino effect of constant recruitment and selection activity. So too, the upsurge of contingent labour (contract and short term contracts) to back fill vacancies gives rise to an enlarged risk profile for the DC organisation. The lack of a pipeline of diverse entrants to the sector diminishes the capacity for a varied landscape of thought, critical analysis that leads to innovation and creative application of technical and

knowledge assets. Couple this with a generally aging workforce and the picture for organisations that are chronically 'unready' is grim.

These few examples of 'big ticket' and often intractable problems exist in a business environment of shifting business models, a sector that is swiftly consolidating and that is enhancing operations in emerging markets. Place all this against a backdrop of more savvy but possibly shrinking operational budgets and a highly competitive landscape it becomes obvious that the business responses of yesterday are little match for the sectoral challenges of today, let alone tomorrow.

When we think of research and development, innovation and advancement we tend to think about the kit; the fancy new widget that will help us make better decisions, build better tools or store and utilise data more securely and efficiently. We rarely think of

the innovations in the way we lead, manage, educate and engage our people within the organisations charged with building the next big thing.

### But we should

Research currently being undertaken by Anglia Ruskin University suggests that, like many other sectors, the DC sector is hampered by traditional business thinking. Inflexible organisational silos are obstructing innovation in the ways companies are organised and traditional 'back end' structures provide incompatible support to a 'front end' that is dynamic and geared for change and challenge.

For example, the ways in which people are recruited, trained, retained and managed in general is largely still influenced by and designed around business practices of mid last century. The attitudes and practices of recruitment and selection for example are still largely predicated on the workforce

structures of that time too; but that time has passed and the workforce itself is remarkably different in structure, behaviour and expectation.

The Data Centre sector, of course, lies at the confluence of the ongoing digital revolution. The 'second machine age' demands a new managerial paradigm. The intersection of rapid and discontinuous technological advances and latent demand for new business approaches is quite clearly the space where the Data Centre sector resides.

It is self-evident then that innovation should not just be thought of in terms of the tech or the gadgets. As the digital revolution marches on unconstrained, so should we see an organisational revolution of leadership capability, managerial expertise and business structure if the 'perfect storm' of sectoral challenges is to be diminished and overcome.

## Aisle containment is not new but has unresolved issues



By Prof Ian F Bitterlin CEng FIET. Consulting Engineer & Visiting Professor, University of Leeds.


WHILST THE PRINCIPLES of thermal management in the data centre, cold-aisle/hot-aisle layout, blanking plates, hold stopping and 'tight' aisle containment are not exactly 'new' it is still surprising to see how many facilities don't go the whole way and apply those principles. Many of those that do seem to be running into problems when they try to adopt a suitable control system for the supply of cooling air.

The principles are simple; if you separate the cold air supply from the hot air exhaust you can maximise the return air temperature and increase the capacity and the efficiency of the cooling plant. This is especially effective if you are located where the external air temperature is cool for much of the year and you can take advantage of that by installing free-cooling coils that cool the return air without a mechanical refrigeration cycle. In this context, the term 'free-cooling' simply translates into 'compressor free' – you still must move air (or water) by fans (or pumps)

so it is certainly not 'free' but it does reduce the cooling energy. The hotter you make your cold-aisle and the colder it is outside (or you use adiabatic or evaporative technology to take advantage of the wet-bulb, rather than dry-bulb) the more free-cooling hours you will achieve per year.

Clearly this doesn't work very effectively when the outside ambient is hot and humid but for most of Europe from mid-France northwards 100% compressor-free operation is perfectly possible if the cold-aisle is taken up into the ASHRAE 'Allowable' region for temperature. Of course, some users/operators don't want to take such risks (real or perceived) with their ICT hardware so limit the energy saving opportunities and some will still fit compressor-based

cooling to be in 'standby'. For some users, the consumption of water in data centre cooling systems, for adiabatic sprays or evaporative pads, is an issue, but it is simple to show that as water consumption increases your energy bill goes down and the cost of water (in most locations in northern Europe) does not materially eat into those energy savings. In fact, if your utility is based on a thermal combustion cycle (such as coal, gas, municipal waste or nuclear) then the overall water consumption (utility plus data centre) goes down as thermal power stations use more water than an adiabatic system per kWh of ICT load. For a rule-of-thumb you can calculate a full adiabatic system in the UK climate as consuming 1000m<sup>3</sup>/MW of cooling load per year, equivalent to about 10 average UK homes. So, what's the problem? To see what



the problem of tight (say <5% leakage between cold and hot aisle, instead of 50%) air management is we should look at the load that needs cooling, e.g. a server. The mechanical designer positions the cooling fans and all the heat-generating components.

The space is small and confined so the model can be highly accurate using CFD as a design tool. He places the 20-30 temperature sensors that are installed in the average server at the key points and someone writes an algorithm that takes all the data and calculates the fan speed to keep the hottest component below its operational limit.

However, we should note the simple model environment – one source of input air at zero pressure differential to the exhaust air and one small physical space with a route for the cooling air. So, the fan speed varies with IT load as it causes the components to heat up. Then we can look at the cabinet level:

- kW load? But which 'kW', 'design', average or even a set of variables between servers?
- Temperature rise? The delta-T varies by hardware manufacturer and design
- Air volume throughput (m3/s) from each server in each cabinet, variable, not fixed
- Position in the cabinet. Top, bottom, middle? Most 42U cabinets in enterprise and collocation facilities are only filled 50-60%
- Missing blanking plates (a common and far too frequent event) and holes not stopped producing some bypass-air

And what to load each server with? Any model is now incapable of representing reality: If the user enables the thermal management firmware on their hardware (admittedly not as common as it should be) the loads' fan speed is controlled by the IT load (or more specifically by the hottest component) then each server will be ramping its fans up and down on a continuous basis as the work flows in and out of the facility.

In a large enterprise facility, the variations of air-flow are totally random. The fan speed variation with ICT load is unique to each server make/model with some varying kW



load from (as good as) 23% at idle to (as bad as) nearly 80% at idle.

This leads us to consider 'how?' to control the air delivery: Often cooling systems are designed to the set-point feed temperature or exhaust temperature or to operate with a slight pressure differential between cold and hot side – although this is most undesirable as the ICT hardware dictates the air-flow demand and its on-board fans are rated for zero pressure differential between inlet and exhaust. That results in air being 'pushed' through the hardware as the server fans slow down. Don't forget that the cooling air will find the shortest path, not the most effective, and the clear majority of excess bypass air is through the load.

However, for air-cooled servers, the whole concept of a small set of very large supply fans (operating together) feeding a very large set of very small fans (servers all operating independently) with no control between them is a mechanical engineering problem. If we add to that a set of very large fans acting to remove/scavenge the exhaust air we have three fans in series and, without a degree of bypass-air, it becomes a mechanical engineering nightmare.

Then we might as well take the opportunity to look at the generally held view that server fans ramp up at 27°C and any temperature higher than that only serves to cancel out any energy saving in the cooling system.

The principle is true but the knee-point is nearer 32°C. The main reason for the 'Recommended' higher limit of 27°C in the ASHRAE Thermal Guidelines is to limit the fan noise from multiple deployments to avoid H&S issues needing the facility to be classified as a noise hazard area. We can now see that, somehow, the cooling system must:

- Present cooling air into the cold-aisle at whatever temperature the users' appetite for risk or energy plan dictates and ignore the return air temperature (you can't control it so why bother to measure it)
- Contain the cold-aisle and use open grates in the raised floor (if used)
- Ensure blanking plates and hole-stopping is complete and remove any servers that are switched 'off' or seal their air-intakes
- Minimise bypass-air but not try to eliminate it as that will make the controls more difficult
- Only supply enough air into the plenum/cold-aisles so that there is the smallest possible pressure drop across the servers – so variable-speed fans are essential
- If you have access to the ICT hardware then you can monitor the inlet temperature at every server, if not, fit temperature sensors in the cold-aisle

Whatever you do don't succeed in pressurising the cold-aisle as there will be risks of internal hot and cold spots inside your ICT hardware with unknown consequences.

Often cooling systems are designed to the set-point feed temperature or exhaust temperature or to operate with a slight pressure differential between cold and hot side – although this is most undesirable as the ICT hardware dictates the air-flow demand and its on-board fans are rated for zero pressure differential between inlet and exhaust

# The GDPR will bring a major sea change to the EU digital assets hosting community



By Garry Connolly, Founder and President, Host in Ireland

THE EUROPEAN DATA hosting industry is on an upswing as the region continues to attract impressive infrastructure investments from a number of content and technology behemoths, including Google, Apple, Microsoft and Facebook. This rapid growth is not only felt at the top, however, as a host of small and medium-sized enterprises (SMEs) are also recognising and capitalising on the growth of colocation opportunities and increased availability of transatlantic bandwidth capacity due to additional subsea cable deployments.

Against this backdrop, a major sea change looms over the EU data hosting community that will fundamentally impact all industry stakeholders, including service providers and data centres, as well as hyperscales and SMEs alike. Following four long years of debate and preparation, the EU General Data Protection Regulation (GDPR) was approved by the EU Parliament last spring and will go into effect on May 25, 2018.

## State of confusion and non-compliance, says survey

The GDPR is making waves throughout the European data industry, as it is set to replace the Data Protection Directive 95/46/EC, a regulation adopted in 1995. Since that year, there hasn't been any major change to the regulatory environment surrounding data storage and consumption, making this the most significant event in data privacy regulation in more than two decades. Unfortunately, this major shift leaves many IT professionals and organisational heads in a state of confusion as they attempt to understand and comply with coming policies.

According to a report by Netskope, out of 500 businesses surveyed, only one in five IT professionals in medium and large businesses felt that they would comply with upcoming regulations, including the GDPR. In fact, 21 percent of respondents mistakenly assumed their cloud providers would handle all compliance obligations on their behalf.



## Enter GDPR, Exit Britain

It's clear that there are many questions surrounding the coming changes to the regulatory landscape, among these, what will be the impact of Brexit and how this new regulation will affect UK-based organisations. Following this decision, the UK government stated its intention to give notice to leave the EU by the end of March 2017. Under the Treaty of the European Union, after giving notice there will be an intense period of preparation and negotiation between the UK and EU with respect to the terms of withdrawal. Once these terms are agreed upon or two years have passed since the original notice, the UK will be officially separated from the European Union. That means that there is a possibility that UK will remain a part of the EU through March of 2019 – roughly a full year after the GDPR takes effect.

This leaves much confusion for British organisations wishing to remain compliant with their government's data privacy regulations. Given the fact that the UK's membership in the EU is likely to linger into 2019, UK-based organisations must take appropriate measures to comply with the GDPR between May 2018 and the official departure date from the EU, at which point they are to conform to their country's

individual legislation which may or may not differ.

Some of the more formidable obligations of the GDPR include increased territorial scope, stricter rules for data processing consent, and heavy penalties for non-compliance. The GDPR also puts major emphasis on the privacy and protective rights of data subjects, whether individuals or corporate entities. These new rules will not only apply to organisations residing in the EU, but also to organisations responsible for the data processing of any EU resident regardless of the company's physical location. Fines for non-compliance are severe, as any data controller in breach of new GDPR policies will receive fines of up to four percent of total annual global turnover or €20 million, whichever is deemed larger.

May 2018 may seem a far way off, however it's important to remain current and educated on the GDPR to ensure your company is prepared for what's ahead. For data centres, arguably the community that will be most significantly affected by this law, the Data Centre Alliance will keep readers abreast of developments regarding this critical issue in accordance with its mission to promote awareness of industry best practices.

# Airflow matters: (Here's why)



By John Taylor, Vigilent Managing Director, EMEA.

IN A PERFECT WORLD, data centres operate exactly as designed. Cold air produced by cooling systems will be delivered directly to the air inlets of IT Load with the correct volume and temperature to meet the SLA.

In reality, neither of these scenarios occur. The disconnect is airflow – invisible, dynamic, and often counterintuitive; airflow is completely susceptible to even minor changes in facility infrastructure. Properly managed, correct airflow can deliver significant energy reductions and a hotspot-free facility. Add dynamic control, and you can ensure that the optimum amount of cooling is delivered precisely where it's needed, even as the facility evolves over time.

## Airflow challenges

Why does airflow go awry, even within meticulously maintained facilities? The biggest factor is that you can't see what is happening. Hotspots are notoriously difficult to diagnose, as their root cause is rarely obvious. Fixing one area can cause temperature problems to pop up in another. When a hot spot is identified, a common first instinct is to bring more air to the location, usually by adding or opening a perforated floor tile.

Temperatures may actually fall at that location. But what happens in other locations? As floor tiles proliferate, a greater volume of air, and so more fan energy, is required to meet SLA commitments at the inlets of all equipment.

Consider that fans don't cool equipment. Fans distribute cold air, while adding the heat of their motors to the total heat load in the room. As the fan speed increases, power consumption grows with the cube of their rotation. Opening holes in the floor to address hotspots is ultimately self-limiting.

Increased fan speed also increases air mixing, disrupting the return of hot air back to cooling equipment and compromising efficiency. We often see examples of the Venturi effect, where conditioned air blows past the inlets of IT equipment, leaving IT equipment starving for cooling.

Ideally, conditioned air should move through IT equipment to remove heat before returning directly to the air conditioning unit. If air is going elsewhere, or never flows through IT equipment, efficiency is compromised.

Poorly functioning or improperly configured cooling equipment can also affect airflow. Even equipment that has been regularly maintained, and appears on inspection to be working, may in reality be performing so poorly that it actually produces heat. And facility managers don't realize when this occurs.

The data is typically not available. Containment is often deployed to gain efficiency by separating hot and cold air. Unfortunately, most containment isn't properly configured and can work against this objective. Even small openings in containment or non-uniformly

distributed load can lead to hot spots. Where pressure control is used, small gaps lead to higher fan speeds in order to properly condition the contained space.

## Better airflow management through dynamic control

So how can airflow be better managed? First you need data. And not just a temperature sensor on the wall, or return air temperature sensors in your cooling equipment. Airflow is best managed at the point of delivery – where the conditioned air enters the IT equipment. Since airflow distribution is uneven, sensing and presenting the temperatures at many locations within each technical room will provide the best visibility into airflow. And instrumentation of cooling equipment can reveal which units are working properly and which machines may require maintenance.



Next, you need intelligent software that can measure how the output of each individual cooling unit influences temperatures at air inlets across the entire room. When a cooling unit turns on or off, temperatures change throughout the room. It's possible to track and correlate changes in rack inlet temperatures with individual cooling units to create a real-time empirical model of how air moves through a particular facility at any moment.

And finally, you need automatic control. Cooling equipment and fan speeds that are adjusted dynamically and in real-time will deliver the right amount of conditioned air to each location. Machine learning techniques ensure that cooling unit influences are kept up to date over time.