Low Carbon Cooling in Data Centres

Julie's Bicycle

Luke Ramsay, Environmental Sustainability Manager

January 2016

Table of Contents

1. Introduction

1.1 Julie's Bicycle

Julie's Bicycle is the leading global charity bridging the gap between environmental sustainability and the creative industries. Our aim is a creative community with sustainability at its heart and our goal is to provide the inspiration, expertise and resources to make that happen.

Julie's Bicycle has an unmatched track record of research specific to arts and cultural activity, which underpins everything we do. Our team brings together environmental expertise, and experience of the arts and cultural sectors and our website constitutes the most comprehensive resource library developed specifically for the arts and culture anywhere in the world.

We work with over 1,000 cultural organisations across the UK and internationally, to help them measure, manage and reduce their environmental impacts. Over 2,000 companies, large and small use the Creative IG Tools, our suite of carbon calculators and our certification scheme is the recognised benchmark for sustainability achievement within the creative industries.

We believe the creative community are uniquely placed to lead and transform conversation around Environmental Sustainability and translate it into action.

1.2 Ken Dale Travel Bursary

The Ken Dale Travel Bursary offers young building services engineers the opportunity to experience technical, economic, environmental, social and political conditions in another country and to examine how these factors impact the practice of building services engineering.

The Ken Dale Travel Bursary has been established by CIBSE to commemorate Ken Dale's contribution to the Institution and the building services profession. The Ken Dale Travel Bursary makes awards available of between £1,500 and £4,000 to CIBSE members in the developmental stage of their career who wish to spend three to four weeks outside their own country researching aspects connected to their field of work and which will benefit CIBSE, their employer, their clients and the profession. CIBSE is especially keen to encourage applicants to take up the award for research that articulates with CIBSE's concern for the environment.

1.3 Rationale and Aims

Many of Julie's Bicycle clients are heavily dependent on activities which use data centres. The use of cloud computing and streaming services in the music and cultural sector are manifold. Some organisations in this sector might operate a virtual server for their website or use cloud storage in their business. On the other end of the scale some organisations will occupy large parts of data centres, operate their own, or have a big influence on how people consume 'culture' and music (e.g. music streaming services). Either way, these sectors will play a significant part in influencing carbon emissions associated with data centres, which currently stand at c.2% of global emissions and are predicted to rise further in the future.

'Digital' has long been a key theme of research for Julie's Bicycle. Building on this theme, this research aims to understand how the energy intensity of data centres can be reduced, providing insights to help inform decisions made across different levels, from operational to strategic. The aims of this report are: -

1. To explain digital infrastructure, its energy and carbon impacts and how the cultural and music sectors are contributing to this.
2. To explain the main types of datacentres, how they typically operate and how their energy performance is measured.
3. To research and understand how the design and operation of datacentres can reduce their energy intensity.

2. Methodology

2.1 Research objectives and methodology

The overarching aim for the research trip was to understand how the design and operation of datacentres can reduce their energy intensity. In developing a methodology for this research it was clear that although guidance and solutions for reducing the energy intensity of datacentres existed, it wasn't always being adopted.

Following further research and conversations with experts in the field I learnt that dogmatic decision making and concerns about security and uptime often meant energy efficiency and carbon reduction weren't always the highest priority. Therefore this research aimed to visit facilities that had prioritised energy efficiency successfully, to learn from their approach and to champion this best practice through case studies.

Research Questions

For different climates and locations:

1. What designs and operation can be put in place to increase the Power Usage Effectiveness of a data centre?
2. How realistic is adoption of the ASHRAE 'Allowable' envelopes (A1 – A4)?
3. What are the barriers for adoption of broader ASHRAE envelopes?

Methodology

- Visit datacentres located in different climates – wet heat, dry heat, temperate.
- Understand cooling strategies employed in these datacentres, and corresponding annualised PUE.
- Objectively assess barriers to adoption of broader ASHRAE parameters using questionnaire to interview datacentre operators.
- All contributors will have an opportunity to provide feedback to the final report before publication.

2.2 Itinerary

Locations were selected in order to encounter a variety of different climatic conditions. Most of the trips were focused in North America as this provides a high level of diversity when it comes to facilities and climate. In addition, Singapore was selected for a site visit because of its tropical climate.

During two weeks in August and a week in September 2015 I had meetings and/or conducted site visits at the following locations:

1. National Renewable Energy Lab. Golden, Colorado.
2. National Snow and Ice Data Center. Boulder, Colorado.
3. Rocky Mountain Institute. Boulder, Colorado.
4. eBay. Phoenix, Arizona.
5. Department of Energy seminar on data centres. Phoenix, Arizona.
6. Equinix data center. San Jose, California.
7. Lawrence Berkeley Lab. Oakland, California
8. Fred Hutchinson Cancer Research Centre. Seattle, Washington.
9. Equinix data centre, Singapore.

Carbon footprint analysis and the corresponding carbon offsetting can be found in the appendix of this report.

3. Our digital world

3.1 Digital infrastructure

Digital technology has become intrinsic to huge parts of our personal and professional lives in the 21st Century and its adoption continues to spread further, from the The Internet of Things and smart cities, to the almost ubiquitous smart phone. Our digital lives are all inevitably underpinned by the internet, which allows us to efficiently store, process and transfer information. The internet can at times seem intangible, but there's huge amounts of physical infrastructure behind our tweets, cloud storage, YouTube cat videos and streamed Spotify playlists.

Figure 1 illustrates how much of our digital technology interacts. Energy is consumed at each of the four points within this digital infrastructure (Device, Access, Network and Server). Whilst our devices can act independently, they spend a large proportion of their time sending and retrieving data over the internet. From sending an email, to a 'Like' on Facebook, a browse of eBay, to saving your photos to Dropbox – each of these activities require the use of a server somewhere in the world (located in data centres) to store, process and retrieve data. Our global appetite for internet data is increasing, driven especially by streaming content.

Figure 1: Simplified schematic of digital infrastructure

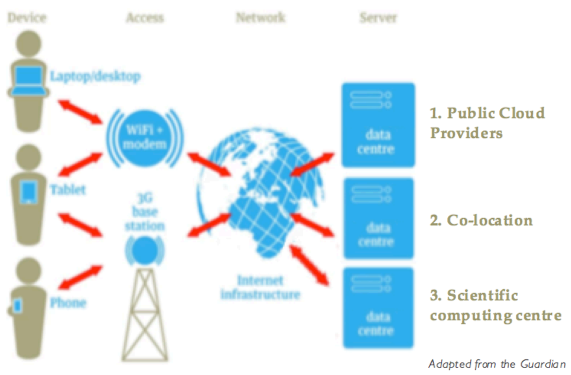

3.2 Internet data

According to research by Cisco[1], global traffic - from datacentre to user - has increased fivefold in the last five years with total traffic predicted to surpass the zettabyte mark in 2016 – that's the equivalent to 2 billion years worth of music![2] The figure is almost four times higher when you include data that either stays in data centres, or travels between them. The report provides current and projected internet traffic for four main sub-sectors: - Internet video; Web, email and data; File Sharing; and Online gaming (figure 2). Between now and 2019 file sharing will remain flat, online gaming will increase fivefold, but still occupy a tiny proportion of internet traffic (0.1%). However, Web, email and data and Internet video are the two subsectors which will drive the biggest increase in traffic, increasing three- and four-fold over this period. By 2019 streaming from Internet video will represent 80% of all online internet traffic, at over 89,000 petabytes per month. A single petabyte is the equivalent to 268 million mp3 tracks.

Figure 2: monthly internet traffic projections, by sub-sector[1]. For reference, 1 petabyte is the equivalent to 268 million mp3 tracks

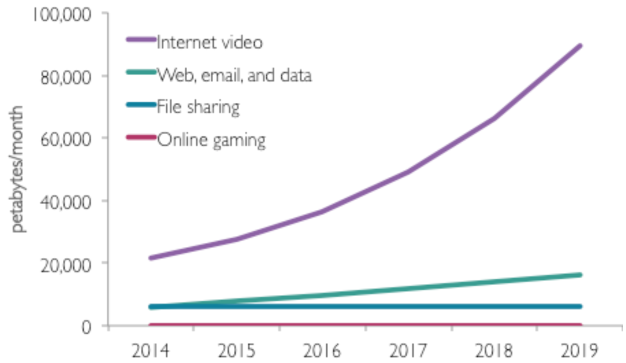

How might the music and cultural sector (see glossary for definition) contribute to this trend? We know that the music sector drives file sharing (typically illegal peer to peer downloading) and that it also drives data, predominately through music streaming and downloading. For example, Spotify - the dominant music streaming service[3] - had 75 million users[4] in June 2015. This shift in consuming habits from physical, to downloaded file, to streaming will continue especially as other giants such as Apple enter the market (Apple's Music streaming service went live in June 2015) - a recent study in the US[5] showed that revenues from streaming alone surpassed that of physical music sales, with music downloads breaking that record long ago, predominantly through the monopoly of iTunes.

Despite this growth, music probably constitutes a relatively small amount of internet data. A recent study by Sandvine[6] showed that iTunes contributed up to 3% of internet traffic in the US and 1.6% in Europe, with no other single music service contributing at least 1% to be included in the analysis. Whilst digital music will continue to grow, its associated data is relatively small when compared to video, for example.

Music videos do constitute a bigger slice of the world's digital pie. Although the music industry has a less direct control on the servers of music videos (the vast majority of videos are viewed on YouTube or Vimeo) this impact should still be acknowledged, especially as it almost certainly exceeds that of digital music. YouTube was shown to occupy almost 20% of internet traffic in 2014 and 27 out of the 30 most watched YouTube videos are music videos[7].

The cultural sector's impacts are less direct and more difficult to scope. Apart from data associated with website hosting and normal business activities (emails, cloud file storage, file transfer), the cultural sector is adopting digital platforms for it's work - from live streaming of performances, to the recent establishment of The Space[8] by Arts Council England and the BBC – the UK's dedicated digital platform for the arts. The cultural sector's contribution to internet traffic is likely to be smaller and more sporadic than that of music.

3.3 Energy and Carbon

Calculating exactly how much energy is consumed at each area of digital infrastructure is complicated. Whilst calculating energy consumed at access, network and server stages is more straightforward, a range of hugely variable factors affect device energy consumption - including the efficiency of the device used and the time spent to process and read a unit of content. Studies quantifying the electricity consumption at these different stages are relatively rare, often because the scope and scale of an inquiry becomes redundant as fast as the internet expands.

A comprehensive study by GeSI[9] (figure 3) provides estimates for carbon emissions associated with ICT (predominantly from energy use) in 2002, 2011 and forecast for 2020. Although growth is declining, ICT will account for a bigger proportion of global emissions in 2020 than it does now. Impacts from datacentres are forecast to increase the fastest – by 7.1% per year between 2011 and 2020, compared to 4.6% for voice networks and 2.3% for end-user devices. Research by Greenpeace[10] shows that data centres and associated networks will consume 1,963 billion kWh of electricity by 2020 – the equivalent to current electricity consumption of France, Germany and Brazil combined.

Figure 3: Global ICT Emissions (GtCO2e)

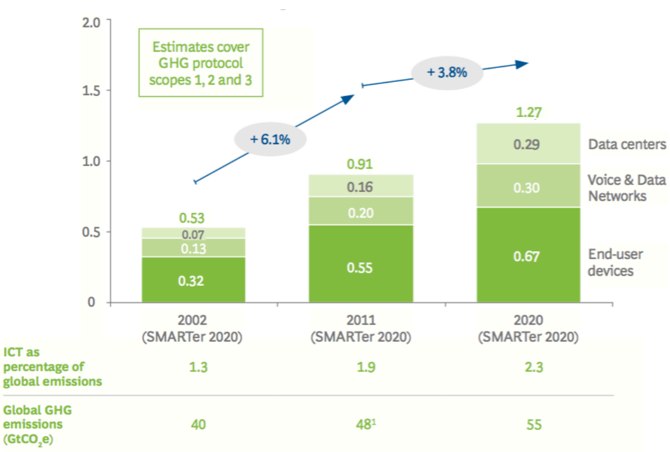

A report published in 2011[11] showed that electricity use from data centres increased by 56% from 2005 to 2010 and was equal to around 1.2% of global energy consumption.

Can our expanding digital infrastructure be a good thing for carbon emissions? The GeSI researchers think so – they identified a range of applications from smart building management to video conferencing that could lead to a 16.5% reduction in the global carbon emissions by 2020 - a saving four times larger than ICT's own carbon footprint. However, whilst ICT might be intrinsic to many of the solutions for a low carbon future, it's important to understand their direct carbon impacts so that we can work to make their operation and use as efficient as possible.

The devices we use on a daily basis are responsible for the biggest proportion of carbon emissions – 60% in 2011. Figure 4 shows the proportion of carbon emissions produced for different device types. For devices that consume relatively low levels of energy during their life time (such as Smart Phones and Laptops) the majority of the carbon emissions are actually created during the production stage of their lives.

Figure 4: Whole life cycle carbon emissions, by device[12]

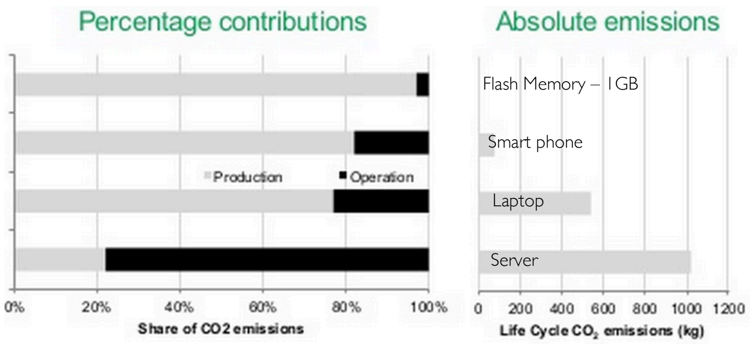

The ratio of operational impact versus production impact is roughly similar for phones and laptops, with around 80% of their life cycle carbon emissions being created at the production stage, often referred to as 'embodied' emissions. For conventional PC's (denoted as 'server') the proportion is the opposite – the majority of energy is consumed during operation.

Average life-cycle emissions for Smart Phones equal 80kg CO2e, laptops come in at over half a tonne and PC's are responsible for a tonne of CO2e over their life time. To put this into perspective, the carbon emissions associated with average UK energy consumption equal c.5 tonnes CO2e. It's estimated that annual global smartphone sales in 2016 will equal 1.4bn[13]; desktop PC's will equal 129, Laptops 202 and Tablets 375 million[14]. Growth in overall sales is forecast at c.2% pa, with tablets and smartphones representing a bigger proportion of devices as users move away from PC's and traditional laptops[15].

Despite PC sales being comparably low, they still comprise 60% of all device related carbon emissions[16], mainly because of their relatively high level of embodied carbon emissions. Therefore our forecast shift away from PC's to devices with lower footprints, such as tablets and laptops will help to reduce the growth rate in overall device-related carbon emissions.

3.4 Atoms Versus Bits

How do our digital activities compare against our more 'physical' ones? As with many of the calculations described in this section, it can often be affected by a range of factors: - type of device, how much time we take to consume a unit of media, how frequently we do that and how else we use that device whilst it's active. However, most studies show that our digital operations are economically more efficient, and often more carbon efficient too. On balance it's better to move bits, rather than atoms!

An example of this - applicable to the music industry – is of physical CD's, purchased from a high-street shop compared with a downloaded piece of music. The results of a study on the life cycle analysis of these scenarios, and some in between, show that reducing the amount of physical infrastructure involved in the process does reduce the associated level of carbon emissions.

Large savings in carbon emissions are made from reductions in transport (customer to shop, CD along supply chain to customers house) and in emissions associated with buildings and warehouses required to store physical stock, when we move towards more digital consuming habits. This study does not include the embodied and energy emissions associated with the devices we listen to the music on, but these impacts would likely be similar across all scenarios.

As shown in Figure 5, choosing a digital download can reduce carbon emissions by 80% when compared to ordering the same piece of music as a physical CD, and the reduction is even bigger when compared to a purchase from a high street shop. 'Burning' your digital copy onto a CD at home, and using a plastic case also results in a reduction of emissions when compared to an online order of a CD - albeit smaller at 40%.
Research by Julie's Bicycle has shown that using a cardboard case can result in emissions reductions of up to 95%[17] when compared to a standard plastic case, but even with this option this research still shows that digital options for consuming music would lead to a lower carbon impact. What's clear from this study – and from the study in figure 3 - is that for our digital activities, data centres (shown in purple in figure 5) constitute a large proportion of carbon emissions. This proportion is likely to become more significant as our devices become more efficient (both in production and operation), whilst at the same time we stream more data intensive media.

Figure 5: Emissions associated with downloading CD's v physical[18]

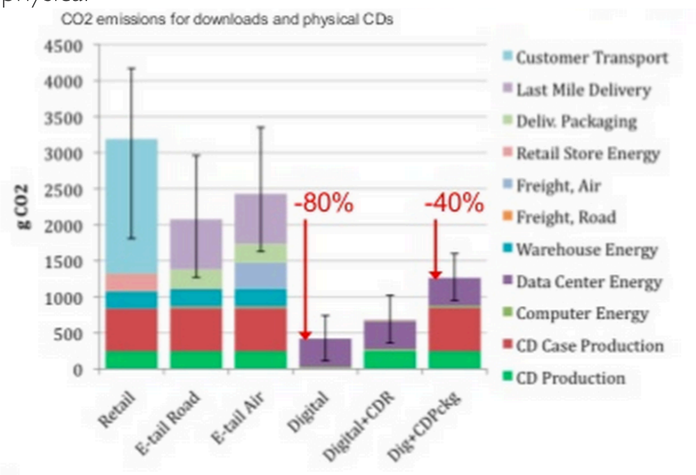

4. Datacentres

4.1 What is a datacentre?

Put simply, a datacentre is a building which houses computing facilities, such as servers, routers, switches or firewalls. Generally speaking, data centres should provide a secure and conditioned environment for this electronic and data processing equipment. The schematic in figure 6 shows some of the necessary engineering required to meet these requirements, including PDU (power distribution unit), CRAC (computer room air conditioning) and UPS (uninterruptable power supply).

Figure 6: Data centre components (CIBSE: Datacentres[19]

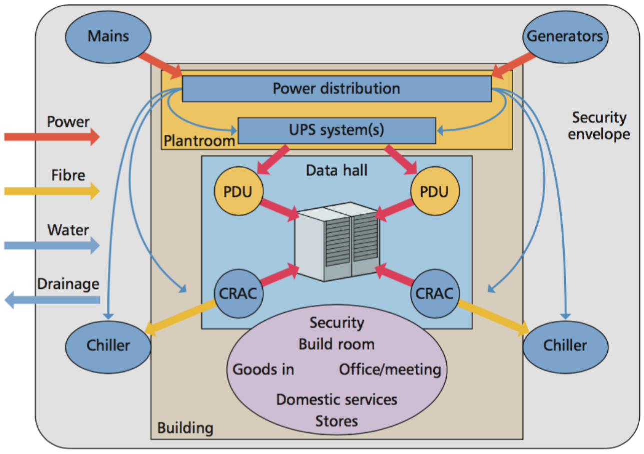

Datacentres are often designed with high levels of redundancy to ensure minimal risk of 'downtime' - i.e. any failures in power or air conditioning that could temporarily reduce or stop the performance of the equipment housed in the facility. Redundancy can be classified in four tiers, based on a methodology developed by the Uptime Institute:

- Tier 1: Non-redundant capacity components. Single uplink. Guaranteed 99.671% availability
- Tier 2: Tier 1 + Redundant capacity components. Guaranteed 99.741% availability
- Tier 3: Tier 2 + Dual-powered equipment and multiple uplinks. Guaranteed 99.982% availability
- Tier 4: Tier 3 + all components fully fault-tolerant including uplinks, storage, chillers, HVAC systems, servers etc. Everything is dual powered. Guaranteed 99.995% availability.

Datacentre tier is often dictated by the nature of its work. For example, a military facility or banking institution might demand a higher Tier than that of a small scale webhosting service. Higher tiers bring more security and reliability, but at significantly higher capital and operational costs. There are four main types of data centres, which each have specific characteristics.

1. Public cloud providers

Public cloud data centres (sometimes called enterprise data centres) typically house servers which provide third party, out-sourced computing. Examples include Google (e.g. servers to store clients email account and Google Drive data and to power search requests), Amazon (Amazon are one of the worlds biggest cloud storage and web host providers) or Rackspace (one of the worlds biggest providers of virtual servers for website and application hosting). Some of the bigger public cloud providers, such as Google, eBay, Apple and Amazon have a sufficiently large enough market and demand that they commission and operate their own bespoke data centres, often with customised servers too.

This homogeneity and greater control of servers means the demand on building services are often less, because air conditioning demands are more consistent and predictable. What's more, an extremely high level of redundancy isn't always necessary for public cloud applications because of the nature of the tasks they perform. This has allowed many public cloud providers to innovate their cooling strategies, by using evaporative and free cooling, for example. For this reason some of the most energy efficient data centres currently in existence are from public cloud providers.

2. Co-location centres

Unlike public cloud providers, that provide a service using their servers, co-location centres typically sell facility space, in which companies house their own servers. Customers of co-location data centres might range from the owner of a single server, right through to a large cloud service provider who might occupy the majority of the centres capacity. The crucial difference with co-location centres compared with public cloud, is that they're selling reliability and security. Customers take responsibility for their own servers whilst the centre is responsible for everything else included in figure 6.

From a building services perspective this diversity of customers (and therefore servers), combined with a service offering based on reliability and security, means that building service design often has a higher level of redundancy and less flexibility on the thermal parameters of the space. This means it can often be more challenging for co-location centres to achieve a high levels of energy efficiency.

3. Scientific computing centres

Typically, scientific computing centres are often referred to as 'super-computers'. These data centres usually house specialised computer systems designed specifically to conduct computationally intensive activities for science and engineering applications from climate modelling to aerodynamic simulations.

These facilities are often very bespoke with the highest power densities of any data centre type. Like public cloud centres, the homogeneity of the servers housed within these centres, and the relatively low uptime demand means there is more flexibility for building design, with consequential increases in energy efficiency.

4. 'In-house' data centres

In-house data centres are those based within the building or campus of an organisation. Traditionally, many businesses have used in-house data centres – ranging from a server cabinet to dedicated room or hall – to serve their in-house IT systems and for web-hosting too. However, as network infrastructures have improved many businesses are moving away from in-house data centres and moving towards cloud based systems. Moving from in-house to public clouds offer much more scalability, higher rates of uptime and reliable support services, not to mention an immediate reduction in on site energy demand for an office. Research shows this shift towards consolidated, cloud based servers results in a net reduction in global energy demand too.

4.2 Metrics

There is a range of standardised energy and sustainability metrics used to account for different types of efficiencies in data centres. The most frequently used is Power Usage Effectiveness (PUE), developed by industry body, The Green Grid. It's defined as the ratio of total facilities energy to IT equipment energy, as shown below.

$$PUE = \frac{Total\ Facility\ Energy}{IT\ Equipment\ Energy}$$

PUE is annualised where possible and is always >1. The global average for datacentre PUE currently stands at around 1.8. Figure 7 shows how energy use breaks down for a facility with a high PUE of 2. In this example, a large demand in energy from compressors (from the use of traditional chillers) has significantly increased the PUE.

Figure 7: Breakdown of indicative energy use, PUE = 2

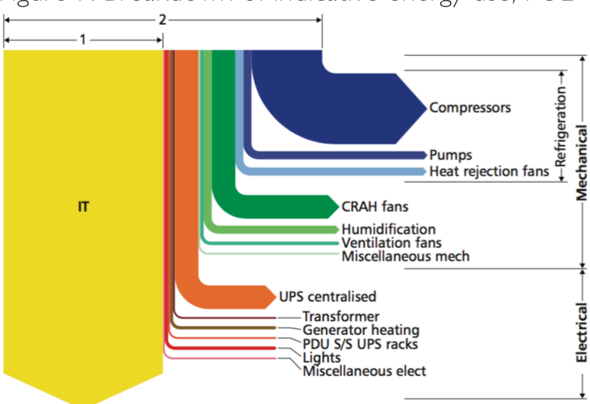

19

PUE has become the dominant metric used in data centres across the world, with the US department of Energy, European Union and the Japanese Ministry of Economy, Trade and Industry agreeing to use the metric in their guiding principles for data centres.

PUE has been an effective tool, increasing transparency and driving reductions in energy use. However, there are still some criticisms of PUE: - despite comprehensive advice provided by The Green Grid, organisations are still interpreting the methodology selectively in order to create the best PUE result for their facility (by missing out sections of infrastructure or discounting for heat recovery, for example); The Green Grid have advised that PUE isn't always comparable between different facilities; and that PUE fails to account for the life cycle emissions of the building materials and servers, some of which can be changed as frequently as every 18 months.

PUE isn't the only metric published by The Green Grid. Other metrics have been developed which are variations to PUE, or concentrate on other impacts such as carbon intensity and water use. They include:

CUE – Carbon Usage Effectiveness. Similar to PUE, but takes account of the carbon intensity of electricity supply. A CUE of 0 would result from sourcing 100% of power from zero carbon energy.

$$CUE = PUE \times Carbon\ Emission\ Factor$$

Where Carbon Emission Factor is carbon intensity of energy sourced for data centre. The units of CUE are kilograms of CO2e per kWh.

WUE – Water Usage Effectiveness. Developed to raise awareness of water consumption from humidification and cooling infrastructure – an increasingly significant issue as data centres move towards evaporative cooling strategies for energy reductions.

$$WUE = \frac{Annual\ water\ usage}{IT\ Equipment\ Energy}$$

The units of WUE are litres per kilo-watt hour.

ERE – Energy Reuse Effectiveness. Introduced to account for the benefit of reusing energy *outside* of the facility, e.g. waste heat used for heating demand in adjacent office. ERE helps stop the misinterpretation/misuse of PUE.

$$ERE = \frac{Total\ Facility\ Energy - Reuse}{IT\ Equipment\ Energy}$$

As with PUE, ERE doesn't have a unit, but mathematically it can be <1 and could almost equal 0 (almost all energy reused outside).

These range of metrics help to evaluate the broader sustainability of a data centre by focusing on different impacts. Like PUE, they should be taken in context of climate (a datacentre with a low PUE in a cool climate will probably be less efficient than a data centre with the same PUE in a topical climate) and other prevailing factors such as data centre type (as in 3.1, type can have affect on level of redundancy and therefore efficiency).

These metrics should also be taken in context of each other. For example a good CUE score (close to 0) might hide the fact that the data centre has a poor PUE (2, inefficient), but still achieves a good CUE score because it uses low-carbon energy to power the facility. Therefore using each metric appropriately, and with the additional context of other metrics is advised.

4.3 Environmental Parameters

The servers and other ICT equipment within data centres create heat - and at increasingly high densities. At the same time, servers also require their environmental conditions to remain within certain ranges in order to operate effectively and reduce the risk of server failure. These environmental parameters, and others such as air particulates and gaseous pollution, are set by the capabilities of the server hardware installed in a data centre.

In 2004 ASHRAE published their first set of 'Thermal Guidelines' for data centres. The guidelines represent recommended and 'allowable' parameters for temperature and relative humidity, based on an assessment of the latest tolerances of computer hardware. Since 2004 ASHRAE have published new iterations (2008 and 2011) as hardware has become tolerant to broader ranges of environmental conditions. One of ASHRAE's guiding principles in the publication of this work is to help drive efficiency whilst maintaining data centre reliability.

Figure 8: ASHRAE Environmental classes (parameters measured at inlet of server)

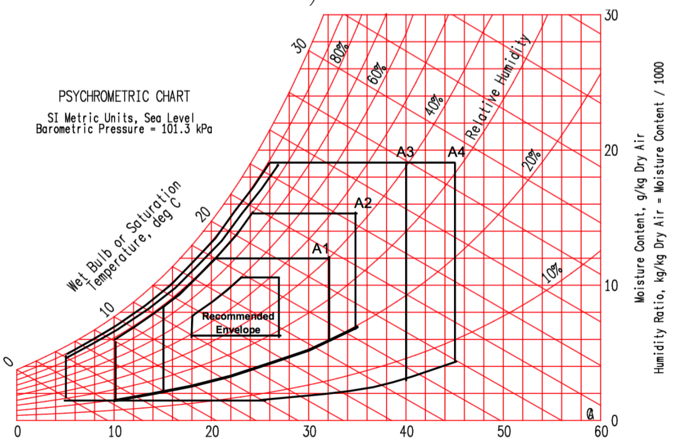

The environmental classes published by ASHRAE are represented in the psychrometric chart in figure 8. The 'Recommended' envelope is intended to be reliable and relatively risk-free, with the proceeding 'Allowable' classes (from A1 to A4) allowing for progressively more flexibility, allowing for less energy intensive cooling strategies (chiller-less, evaporative cooling, free cooling etc.).

However, for the majority of data centres, optimising energy consumption and setting an appropriate environmental operating envelope is complicated and depends on a web of inter-related factors. These all have an affect on the overall PUE and total cost of ownership and the right balance has to be struck depending on the type of datacentre, its location and the resilience/capabilities of the hardware housed within it. Some of the main factors include:

Reliability: In some instances, increased temperatures will degrade the long-term reliability of some components. Choosing more generous operating envelopes in hotter climates has been shown to increase the probability of server failure. ASHRAE recommend using time-at-temperature histograms - defined by the selected operating envelope and using location climate data. This information is then used to calculate a weighted failure risk, based on standardised failure rates at given temperatures.

Server power: As inlet temperatures to server rise, so do fan speeds, to keep components within their temperature limits. Increased temperatures also cause devices to consume more power due to leakage currents. For example, increasing temperatures from 15C to 30C will increase server power by 4 to 8%.

Performance and allowable envelope of server: For a given operating temperature, some systems may perform better than others, depending upon the servers power and thermal design points. A system designed to handle any potential workload compared with one that is tuned for a specific workload may be capable of better overall performance in wider environmental ranges.

Corrosion: Expanding the high end of the humidity envelope increases the opportunity for corrosion problems to occur. Moisture may exacerbate corrosion problems in data centres where corrosive gases are present.

Cost: To support higher inlet tempertures, more costly thermal solutions may be required (such as hot/cold aisles, granular cooling to specific racks, increased monitoring equipement, complex chiller-less cooling strategies etc.). In addition, larger, more expensive heat sinks may be required on servers to maintain acceptable levels of compute performance and power consumption.

5. Site visit case studies

5.1 National Renewable Energy Laboratory – ESIF High Performance Computer

High Performance Computer, ESIF
Location: Golden, Colorado.
Facility type: Supercomputer
Facility size: 16,900m^2, 10MW capacity
Annualised PUE: 1.06
ASHRAE Class: A1
Tier 1

Key technologies – In-rack water cooling, heat recovery, evaporative cooling


Credit: NREL 26387

The Energy Systems Integration Facility (ESIF) High Performance Computer (HPC) is a super computer that has been in operation on the NREL campus since 2014. The computer is used for a wide range of applications associated with research projects taking place at NREL, and also remotely from other research laboratories around the world. It has a peak performance of 1.2 petaflops - this is equivalent to the power of 15,552 laptops! The high temperatures created by the power dense servers (50kW/ rack) make this data centre an appropriate choice for a direct liquid cooling strategy.

Building services:

90% of the heat load from the HPC is extracted from the room using an innovative direct-to-chip cooling system, where copper pipes containing relatively warm water (24°C) are fed across the computer chips and circuit boards. The increased thermal capacity of water (relative to air) enables this relatively high temperature. This higher temperature means cooling demand can be met all year round by using four evaporative cooling towers, relying on no mechanical cooling. The high grade heat is also recovered and used in neighbouring buildings for heating and some hot water demand. The remaining 10% of the heat load is managed through air cooling using hot aisle containment, serviced by a fan wall connected to the same chilled water loop. This strategy has resulted in an incredibly low annualised PUE of 1.06.

Key learnings and insights:

- Direct liquid cooling provides a highly effective means of cooling servers, especially in high power density applications such as super computers. Direct water cooling is also very efficient, since the relatively higher temperatures of the chilled water loop mean low energy cooling can be used - in this case, evaporative cooling towers.
- The relatively high grade heat created from the liquid cooling loop can be used to service local heating demand too. However, compared to conventional heating systems this grade heat is still relatively low, so heating demand within the same building or within a local campus is preferable.
- Water purity is key in liquid cooled systems – algae and paint flecks affected this system.
- Evaporative cooling towers are an efficient form of cooling for this application, but they do have a high water demand. Currently, this data-centre uses 7,600m^3 per MWh of power consumed. There are plans in the near future for the cooling demand to be supplemented with a thermosyphon dry cooling system, which will reduce overall water demand during night-time and in winter months, when temperatures are consistently below the cooling liquid temperature of 24°C.
- 'Legacy' specification computers in a data centre can impose limits on the thermal parameters of the data-centre. For example, older tape-based storage systems were still being used in this data centre, requiring air-cooled ventilation, with stricter humidity controls. Physically enclosing older systems such as these can help control the conditions independently and therefore reduce overall demand in the building.

Further information:

ESIF data center website
http://svlg.org/wp-content/uploads/2013/11/NREL-Data-center-case-study.pdf

5.2 National Snow and Ice Datacenter

National Snow and Ice Datacenter
Location: Boulder, Colorado.
Facility type: Local network datacentre
Facility size: 37.8kW, 0.0378MW
Annualised PUE: 1.25
ASHRAE Class: Allowable class A1

Key technologies – Indirect evaporative cooling, free cooling, photo-voltaic charged UPS.

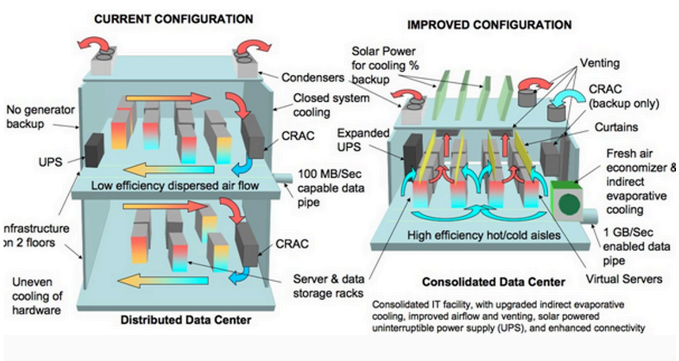

The National Snow and Ice Datacenter (NSIDC) processes 120 terabytes of scientific data used by researchers worldwide. If you've recently read about retreating glaciers or thinning ice caps then the chances are it was based on data from this facility, direct from NASA satellites. The data centre is relatively small (37.8kW) and is located within a single large room of an office building. It provides an interesting case study because it underwent a comprehensive refurbishment in 2011. The refurbishment and new cooling strategy realised an 88% reduction in cooling electricity demand when compared to the previous, more conventional direct expansion cooling system.

Building services:

This datacentre employs a mixed strategy of free cooling and indirect evaporative cooling. For this datacentre, operating at ASHRAE Allowable class A1, free air cooling can be used when outside temperatures are relatively high - circa 30°C. When cooling is required (or incoming air humidity drops below 25% RH), then indirect evaporative cooling is used. The eight indirect evaporative cooling units use an innovative approach which harness the highly efficient Maisotsenko-cycle to cool incoming air to below dew-point temperature, by around 10°C. Being indirect, the cooling system does not add moisture to the cooled 'product' air, as it exhausts the saturated 'working air' to the atmosphere, unless it's needed to add humidity to the room.

Air handling is managed using a hot-aisle containment strategy, enabling higher intake temperatures and therefore reduced energy demand. An air handling unit keeps the room under some pressure, which combined with the rising contained hot air, means there's no need for extract fans.

Overall electricity demand of the datacentre is reduced further with the installation of photovoltaic panels on the roof of the building, which produce 36,000kWh of electricity on average per year. The uninterruptable power supply (battery power storage) for the datacentre is charged using the photovoltaic array, with excess electricity being fed back into the state grid system.

Key learnings and insights:

- Because of the dry climate in Boulder, the evaporative cooling units still need to be activated to provide humidity to the space even on cooler days. Although energy efficient, they do use a relatively high amount of water - the 8 cooling units at this facility are estimated to consume $0.3m^3$ per hour in moderate cooling mode.
- Because the cooling system relies on simple mechanics they have a low maintenance cost - c.$2,000 pa.
- Hot-isle containment is an effective way of reducing overall energy demand. Containing the heated exhaust air from the servers prevents it from mixing with incoming cool air. This means the incoming cool air can be set to a higher temperature set point, reducing energy demand.
- Server virtualisation has helped consolidate the number of servers within the datacentre, increasing the overall computational power per watt of electricity.

Further information and links:
https://nsidc.org/about/green-data-center/project.html
http://blog.rmi.org/blog_making_big_cuts_in_data_center_energy_use
https://nsidc.org/about/green-data-center
http://nsidc.org/about/green-data-center/stats/current.jpg
http://symparchive.uptimeinstitute.com/images/stories/Symposium_2012_PPTs/gallher-osbaugh-dc-retrofit.pdf

5.3 eBay, Arizona

eBay datacentre, Arizona. (Modular system).
Location: Phoenix, Arizona
Facility type: Enterprise data centre, Tier II
Facility size: 1MW per modular container
Lower recorded PUE: Dell Epics - 1.03.

Key technologies – Direct evaporative cooling, free cooling, fully modular design, containerisation.

eBay has a range of data centres across North America, but their facility in Phoenix is exposed to the highest average annual temperatures. In Phoenix, the average annual temperature is 24°C with summer daytime temperatures frequently reaching 42°C and above. The data centre is split on several floors, but this case study focuses on the modular rooftop deployment. The container based modular units here present an interesting innovation, proving the ability to achieve good levels of PUE even in challenging hot climates.

Building services:

These modular data centre units are built within a typical freight container design. The units are pre-fabricated with all building services and servers fully integrated and customised to meet eBay's needs. Two different units have been purchased for this site – the Dell Epic and the HP Performance Optimized Datacenter (POD). Although made by different competitors, the units operate using similar designs. These units were created using an integrated design approach and this more holistic approach has helped realise an energy efficient design. For example, the server racks were designed by the same team as the architect designers. With this approach they realised fans could be removed from the servers (by delivering air to the servers only with AHU's), creating a higher rack density and a relatively lower proportion of energy demand for cooling.

The units have a high rack density (27kW per rack), and are designed with very effective hot/cold isle separation. Overall this reduces relative cooling demand – inlet temperatures can be set higher, and the cooling system requires relatively smaller volumes of air to cool the servers. Depending on outside air temperatures, cooling demand is met using direct evaporative cooling, heat recovery ventilation or free air cooling.

Key learnings and insights:

- Modular units make deployment to site very effective, with a fully integrated and pre-fabricated data centre almost ready to operate from the day of delivery. They can provide very energy efficient data centre solutions, even in particularly hot climates.
- The modularity of the system also means that capacity can easily be increased quickly and without the requirement of further building infrastructure.
- The relatively small air volumes and hot/cold-aisle containment also increase the effectiveness of the cooling system to increase overall efficiency even further.
- However, for a large enterprise data centre organisation modular units may not be the most cost effective way of increasing capacity compared to constructing large halls or warehouses.

Further reading:

Modular Datacentres http://www.ellipticalmedia.com/partnerportal/documents/LawrenceBerkeleyModularDCStudy.pdf

5.4 Equinix co-lo facility, San Jose

Equinix SV5
Location: San Jose, California
Facility type: Co-location
Facility size: 11,900m²
Average annualised PUE: 1.4

Key technologies – Evaporative cooled refrigeration system with refrigerant side economizer. 'Granular' cooling, allowing data hall to be cooled at cage level. LEED and Energy Star certified buildings.

This facility in San Jose is one of seven that Equinix operate across 'Silicon Valley'. Being a greenfield construction the design has benefitted from some key features to help increase the efficiency of the data centre. The data centre has been commissioned in three phases spanning several years.

Building Services:

For phase 1 and 2 cooling demand is met through an evaporative cooled refrigeration system with refrigerant side economizer. Phase 3 uses indirect evaporative cooling, with air side economiser for added efficiency. Air is delivered to different parts of the building using a dedicated cold aisle arrangement. The design of the data centre means cooling demand can be met using a 'granular' approach, where cooled air can be directed to specific server cages using overhead ducts. This has helped realise big reductions in overall energy demand, by avoiding unnecessary cooling of cages where there are few servers, or low loads. Cooling is provided through a tiered system installed in the roof, rather than via an under floor system. All pumps and fans used across the data centre operate using variable speed drives, which also allows the cooling system to modulate for demand.

At the time of writing, Equinix were committed to installing 1 MW of on-site renewable energy generation, from bio-gas fuel cells. 1 MW of clean energy will be provided using Bloom Energy fuel cells. The system is predicted to provide 8.3 million kilowatt-hours per year of clean energy direct to the facility.

Key Learnings and insights:

- In co-location facilities it's very important that a cooling strategy can be dynamically adjusted to meet the demand in different areas of the building at different times. The 'granularity' of the system in SV5 shows how effective this can be by achieving a strong PUE.
- Although the air handling system does allow for a 'granular' approach to cooling across the data hall, it works most effectively when server cages within aisles are fitted to full capacity. Delays in bringing the data centre up to full capacity result in extended periods of energy inefficiency at the centre.
- Customers located in a co-location facility can have a large influence on overall energy efficiency. A key part of Equinix's Service Level Agreement (SLA) is to maintain high levels of uptime and consistency in terms of the thermal parameters within the data centre. The thermal parameters are set by Equinix on a global level, by taking into account the latest ASHRAE recommendations, and the demands set by their customers. Some co-location customers are still driving demand for conservative thermal conditions, which will have an adverse effect on energy consumption (i.e. by demanding low levels of variation in inlet air temperatures and strict humidity variations). Educating customers on the reliability of their equipment in the context of broader temperatures and humidity is key to achieving higher levels of efficiency within data centres. Equinix work hard on this issue to ensure good levels of energy efficiency across their global sites.

Further reading:

http://www.equinix.co.uk/locations/united-states-colocation/silicon-valley-data-centers/sv5/
https://www.energystar.gov/index.cfm?fuseaction=labeled_buildings.showProfile&profile_id=1027648

5.5 Computational Research and Theory Facility, Lawrence Berkeley Lab

Computational Research and Theory (CRT) Facility
Location: Oakland, California
Facility type: Acadmic supercomputer.
Facility size: 7.5MW
Average annualised PUE: Yet to commence service, aiming for PUE of 1.08.
Rack density: 65 – 100kW

Key technologies: Free cooling with wet cooling towers, liquid cooling, heat recovery, hot isle containment.

The Computational Research and Theory Facility (Wang Hall) is Lawrence Berkeley Lab's newly constructed super computer. At the time of writing it was preparing for deployment of one of the worlds most powerful supercomputers. The facility will be used by the National Energy Research Scientific Computing Centre (NERSC) and other academic institutions - comprising 6,000 users from 47 countries around the world - to support research on critical global issues like climate change. Apart from housing one of the world's most powerful super computers, this centre is also located almost directly on top of an active tectonic fault line, presenting an additional set of unique challenges.

Building services:

Cooling demand is met using cooling towers (3,000 tons, with air side and water side economisers) linked to a liquid cooling system for the servers. The liquid cooling system comprises liquid cooled radiators mounted in the server racks, with transverse air cooling to the rack. The liquid system allows for higher inlet temperature and therefore evaporative and free cooling provide sufficient demand all year round without the need for chillers. Apart from higher inlet temperatures this cooling strategy is also possible because of the relatively cool maritime temperatures around San Francisco. 93% of all heat generated by the computers in this centre is rejected through the water cooling system with the remaining 7% of heat rejected to the room. This heat can be recovered and used to serve heating demand in the offices in the floors above the data centre itself.

The CRT has been built practically on top of an active tectonic fault line – the Hayward Fault. This provided some unique challenges to the designers of the building. All of the servers are mounted on an seismic isolation floor, which provides 18 inches of lateral movement in the event of an earthquake.

Key issues and learnings:

- Cooling towers will enable the CRT to realise a very efficient PUE level. However, this building was originally specified when droughts were less prevalent in the state of California. Along with a bigger design budget, consideration of the water footprint of the data centre might have led to an alternative cooling strategy met with dry coolers instead.
- Air and water inlet temperatures for the servers are maintained at around 18°C with a maximum of 24°C. These inlet temperatures are lower than the upper limit recommended by ASHRAE, as it's harder to achieve higher temperatures with the climate around Berkeley. Interestingly, the CRT team calculated that although higher temperatures would reduce the load on services, it would actually increase the computing load (leading to a net increase in power) because the chips use more power for the same computational load when they're hot. This is an issue that matters more for super computers with very high rack densities.
- The data centre has realised such good levels of PUE because it has a cooling strategy that allows for fluctuations in relative humidity. However, this is an issue for the older tape based storage also contained within the data centre. In order to avoid these computers being exposed to large fluctuations in humidity they've been isolated from the rest of the computers within the room. Tapes still represent an energy efficient form of data storage because they can hold large volumes of data in a way that is energy passive.
- PUE can be calculated using different methodologies, which bring different resulting scores. The PUE is very low at CRT because the fans used for the inter-rack cooling are considered to be part of the computer power rather than facilities. If they were considered to be part of the facilities then the resulting PUE with current equipment would be closer to 1.2.
- The CRT will have a very high level of server utilisation (98%) throughout the day and across the year and this is ensured through the use of 'batch scheduling' jobs where between 150 and 200 jobs can be scheduled to run simultaneously on the supercomputer. In the rare event that temperatures become too high for the data centre at full capacity then the chips within servers are automatically turned down using power capping. This approach has been taken instead of relying on mechanical cooling to cool the servers and keep them running at a higher capacity.

5.6 Fred Hutchinson Cancer Research Centre

Fred Hutchinson Cancer Research Centre, Eastlake Data center
Location: Seattle, Washington
Facility type: Local network datacentre, Tier II
Facility size: 750kW
Annualised PUE: 1.1
ASHRAE Class: A1

Key technologies – Direct evaporative cooling, free cooling, mechanical cooling for hot humid days, hot aisle containment.

This data centre was designed and built to accommodate servers used for research taking place at the Fred Hutchinson Cancer Research Centre. The newly built data centre will host a consolidation of a range of servers from across this research campus. Taking advantage of the relatively mild maritime climate of Seattle, the data centre has realised a very low PUE of 1.1.

Building services:

The data centre predominantly uses direct evaporative cooling and free cooling (90% of annual hours), with mechanical cooling (chilled water coils) being used on hotter, more humid days. The building was designed so that intake air comes from a shaded location, reducing the demand on the cooling systems. Designed sequentially, outside ambient air is drawn through four sections: it's first filtered for particulates, then cooled using direct evaporative cooling if necessary, and cooled further via mechanical cooling (chilled water coils) when outside air temperatures are particularly high. As shown in the diagram below cooling is fully modulated using a bank of fans to move air through the different air conditioning stages, with the server room being kept slightly above atmospheric pressure to reduce the demand on server fans.

Resulting hot air from the servers is fully contained in hot aisles, which help reduce overall energy consumption from cooling. This is because incoming air does not mix with outgoing hot air, meaning cooling temperatures can be set higher.

Key learnings and insights:

- Because the data centre began operation before it was at full capacity there was limited airflow through the cooling sections. This caused growth of some mould on the material used in the direct evaporative cooling system.
- Because of the relatively low-tech design of the data centre, the overall build costs were half that of an equivalent fully mechanically cooled datacentre.
- Comprehensive metering across the facility helps identify where hot spots are within the data hall and also provides real time energy consumption of servers and racks. This level of monitoring is crucial to realising an efficient facility – by being able to see the heat distribution across the hall and understand energy demand, the building managers are able to work proactively to solve issues that may result in inefficient operation.

Further reading:

http://www.aeieng.com/index.php/10729-00

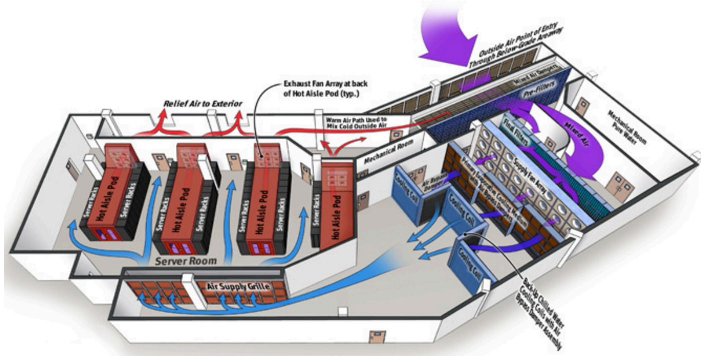

5.7 Equinix co-lo facility, Singapore

Equinix SG3 datacentre, Singapore
Location: Singapore
Facility type: Co-location data centre
PUE: Predicted – 1.4
Facility size: 22 MW, 35,000m2, 7,000 racks
Awards: BCA-IDA Green Mark Platinum Award, SS564:2013

Key technologies – chilled water system with evaporative cooling towers, using grey water harvesting, on-site solar

This data centre is currently Equinix's newest and biggest in Singapore. At the time of visiting the site was in the final stages of commissioning. The data centre draws on experience from across Equinix's other global sites to incorporate a range of building services to realise a high level of power usage effectiveness in the challenging climatic conditions of Singapore. This facility comprises of five floors, with each floor housing between 1,200 and 1,500 racks each. Like all of Equinix's centres across the world this facility acts as a key interchange for many web-based services.

Building services:

Despite being located in Singapore – with a consistently hot and humid climate – this datacentre is predicated to achieve a strong PUE of just below 1.4. This impressive level of PUE has been achieved through careful management of the cooling load based on live data on heat output from across the datacentre hall.

Cooling demand is met through a chilled water system linked to Computer Room Air Conditioners (CRACs), with heat being ejected through evaporative cooling towers. This method of cooling does require a relatively high amount of water, but the system uses rainwater harvesting to reduce the overall potable water footprint. The CRACs are kept on 'hot standby' and work in series to spread the load, which ultimately reduces overall energy consumption. In addition, condenser water pumps are mounted externally, which reduces plant room cooling load. Cold aisle containment has been installed across the site and the halls use raised floors to distribute cold air. The floor is raised to 1.1 m – optimum height to minimise air flow resistance.

Data on cooling demand across the data hall is provided using Synapsense - a wireless monitoring and cooling control solution. The Synapsense system has been integrated into the data halls to provide a very detailed picture of cooling demand at any given time, allowing the cooling plant to be modulated accurately to reduce energy waste. The cooling ventilation within each of the data halls is 'granular', meaning that cooling can be managed using this data across small areas of the hall – adapting to the cooling demand of small areas of racks rather than large whole sections of the room. The system also uses the data to create a 'heat map' within the data hall. This is a really important feature as it can help identify hotspots – areas where obstructions, faulty servers, or inadequate ventilation could cause overheating.

Humidity within the data hall is reduced using desiccant dryers, much more efficient than using chillers for this purpose. Because the data halls are enclosed spaces de-humidification isn't required 24/7 – air is constantly being re-used and re-conditioned.

There's also a small amount of renewable energy installed at the facility – the array of PV panels are rated at 200kW.

Key learnings and insights:

- Effective use of heat sensing data in this facility has helped create a dynamic and precise cooling strategy, minimising unnecessary energy consumption from cooling.
- The design team were originally advised not to fit a variable speed drive to their chillers because of the tropical climate. However, after installation they realised that this could still achieve energy savings. They've now retrofitted chiller optimisation to help increase energy efficiency of this plant.
- Rainwater harvesting has been an effective way of reducing the potable water footprint from the evaporative cooling towers. Another additional benefit of using harvested rainwater is also because the pH is more favourable for this system.
- Cold isle containment was selected because it's more practical for maintenance/client access in a co-location centre. It also makes installing fire extinguishers more practical.
- Hotspots are particularly problematic as cooling systems often cool according to the highest temperature reading in a given space. Using the synapsense system hotspots can be identified proactively, and facility managers can react quickly to avoid unnecessary energy consumption.

6. Insights and conclusions

Apart from the insights raised in each of the case studies in section 5, this research also identified some common themes and issues across the industry.

Changing approaches and attitudes

A key issue is that many decisions made for datacentre operations are made by a different team, often with little consultation with the building management team. This means building managers often have to make compromises, which lead to inefficient operation of facilities.

A holistic approach – including both the building designers and hardware/IT operator within the same budget is an effective way to ensure everyone involved in the design and operation of a datacentre is concerned with energy efficiency. Sometimes known as the 'Chips for Bricks' principle, this approach can consolidate and reduce budgets for data centre projects by linking capex and opex budgets. It also ensures the computer hardware and building services are well matched, making for energy efficient facilities.

Dogmatic approaches and beliefs from a wide range of stakeholders can still hinder innovation. Some of the projects visited in this research had to work through 3 different design teams before they found one that understood the vision for a low carbon facility. All previous teams weren't willing to take the 'risk' of being involved in project using new innovative cooling strategies.

On the end of the relationship are data centre clients – those that use the space to house their servers. Many clients still demand very stringent temperature controls for their servers in the belief it increases reliability. Some even go to the effort of installing thermometers on their hardware to monitor their datacentre. Data centres (co-location ones especially) therefore need to ensure they're always educating their users on thermal parameters, and the capabilities of their servers, as these become broader.

Customers using cloud computing (for web hosting, file storage etc.) also have role to play by informing themselves and choosing providers with strong track records on low-carbon, energy efficient operation.

Responsible resourcing

Focusing on reducing relative levels of energy consumption in data centres has been helped by metrics such as Power Usage Effectiveness (PUE), which have created transparency and healthy competition in the industry.

Going forward, the industry should also focus on increasing efficiency in other key areas such as carbon – making energy source just as important as energy intensity. This research has shown that some datacentres are already making big commitments to buying their energy from renewable sources. As big consumers of energy, datacentres have the opportunity to make a real difference to energy production in the regions they reside by investing in renewable energy.

Water efficiency is also particularly important in draught vulnerable locations of the world where data centres are now have a significant impact on water supplies, especially through the use of evaporative cooling methods. It's also important to consider the 'embodied' impacts of a data centre – including the building materials and all the hardware stored inside.

Monitoring is managing

Having a detailed understanding of how much heat servers are producing - as well as a range of other factors including server inlet temperature, humidity and server power – is fundamental to an efficiently operated data centre. For example 'hot-spots' can be quickly identified – servers which may be faulty, over-heating and causing huge unnecessary demands on the cooling systems.

Heat mapping also helps to identify areas where air management could work more effectively. They build the case for relatively cheap interventions, such as hot aisle/cold aisle separation. Many data centres suffer from hot and incoming cool air mixing when server racks are removed, or even when technicians trolley's are accidentally left near a vent! Better airflow management permits warmer supply temperatures, and ultimately lower energy consumption.

Zombies

It's important to focus on how data centres are operated, but we shouldn't lose sight of the IT perspective. 'Zombie' servers – those that are no longer preforming any work, but continue to operate are a big draw on data centre energy consumption. The uptime institute reports that between 15 and 30% of servers are not being used, with estimates putting it at 10 million servers worldwide – that's the equivalent of power from eight large power plants!

Server utilisation could be vastly improved too, by encouraging more users to operate virtual servers, rather than single servers, which can't react dynamically to changes in demand and overall consume more energy. Ensuring server utilisation rates are high and old servers are checked and shut down will be increasingly important.

7. Appendix

7.1 Carbon emissions and offsetting

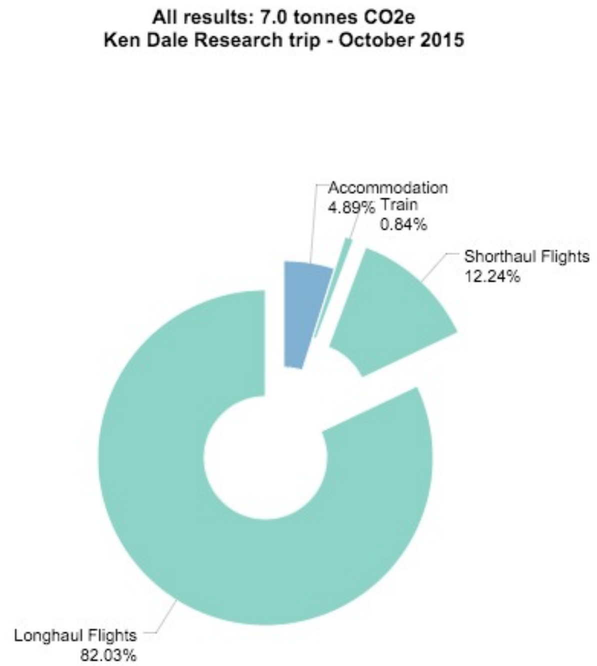

www.ig-tools.com

With locations across North America and South East Asia, this research trip resulted in just over 27,500 miles of travel, predominately by air. Train travel was possible between San Francisco and Seattle (800 miles), but transport infrastructure and limited time meant internal flights had to be taken between some visits.

Carbon emissions were calculated using the International GHG protocol methodology. Carbon emissions associated with all travel and accommodation totalled 7 tonnes CO_2e. This is roughly the equivalent to the average energy emissions of 1.5 UK households. Carbon impacts from accommodation were a third lower than average as apartments (through Airbnb) were selected over hotels for the majority of stays.

These carbon emissions were 'off-set' using sandbag, a UK-based non-profit organisation. They purchase and 'retire' EU Emissions Trading scheme permits, therefore reducing the total carbon budget for emitters affected by the scheme.

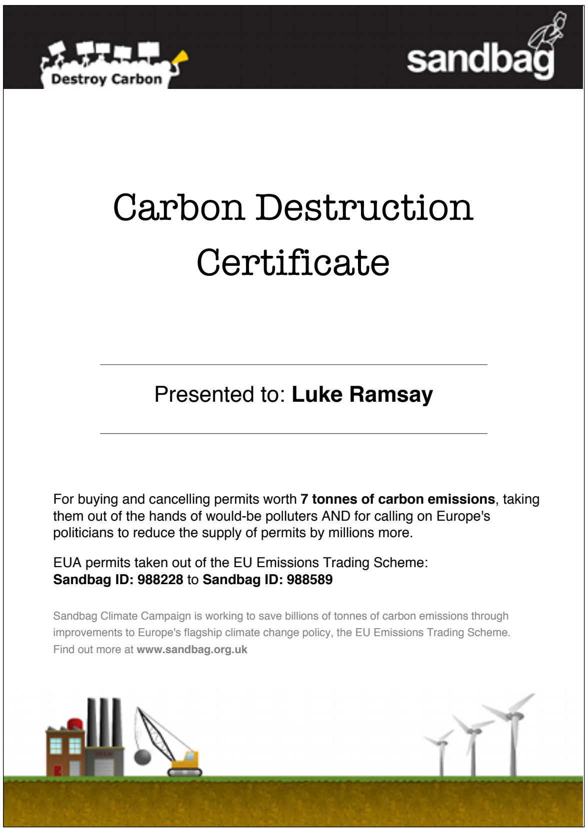

Destroy Carbon

sandbag

Carbon Destruction Certificate

Presented to: **Luke Ramsay**

For buying and cancelling permits worth **7 tonnes of carbon emissions**, taking them out of the hands of would-be polluters AND for calling on Europe's politicians to reduce the supply of permits by millions more.

EUA permits taken out of the EU Emissions Trading Scheme:
Sandbag ID: 988228 to **Sandbag ID: 988589**

Sandbag Climate Campaign is working to save billions of tonnes of carbon emissions through improvements to Europe's flagship climate change policy, the EU Emissions Trading Scheme. Find out more at **www.sandbag.org.uk**

7.2 About the author

Luke Ramsay works as Environmental Sustainability Manager at Julie's Bicycle. For the last four years he's worked with arts venues and festivals across the UK, Europe and Asia, helping them understand and reduce their environmental impacts. Luke heads up Julie's Bicycle Creative Industry Green certification and is responsible for the development of their unique carbon calculators – the Creative Industry Green Tools. He is a qualified Low Carbon Energy Assessor (CIBSE), ESOS Lead Assessor and IEMA Associate member.

7.3 Acknowledgments

I owe a huge debt of gratitude to the following people, whose generous support made this research possible.

- All at CIBSE and those associated with the Ken Dale award, especially Grace Potthurst and Carilyn Clements - thanks for all your help throughout the research.
- Professor Ian Bitterlin – your insights and knowledge were extremely useful in helping to refine my research proposal.
- Emma Fryer, Tech UK – thanks for your insights and for sharing your network, allowing me access to facilities that would have otherwise been impossible to visit.
- Claire Buckley and Alison Tickell, Julie's Bicycle – you were open and supportive from the very beginning, thank you so much for providing me the opportunity to pursue research outside of my day job!
- Otto Van Geet, National Renewable Energy Lab – thanks for being so generous with your knowledge and contacts and for such an informative tour of NREL.
- Dale Sartor, Lawrence Berkeley National Laboratory – thanks for your efforts in ensuring I could take part in the eBay site visit and the data centre workshop in Phoenix.
- Special thanks also go to all others who agreed to participate in the research: - David Gallaher, National Snow and Ice Data Center; Richard Reyher, eBay; Bill Strong, Jim Dawson and Warren Weitz, Equinix San Jose; Craig Schiller, Jonathan Walker, Chris McClurg, Rocky Mountain Institute; Jeff Broughton, NERSC; Jim Walker, Fred Hutchinson Cancer Research Center; Ted Brown, Seattle City Light; George Seet and Sudarmanto Wahyuaji, Equinix Singapore.

7.4 References

[1] CISCO. 2015. *Cisco Visual Networking Index: Forecast and Methodology, 2014-2019 White Paper.*
[2] XO Communications. 2015. "*2016 The Year of the Zettabyte*". http://www.xo.com/resources/infographic/2016-the-year-of-the-zettabyte
[3] Statistica. 2015. *"Spotify Ahead In Terms of Paying Users".* http://www.statista.com/chart/3899/paid-subscribers-of-music-streaming-services/
[4] Statistica. 2015. *"Number of global monthly active Spotify users from July 2012 to June 2015 (in millions)"* http://www.statista.com/statistics/367739/spotify-global-mau/
[5] RIAA. 2015. *2015 Mid-Year RIAA Shipment and Revenue Statistics.*
[6] Sandvine. 2014. *Global Internet Phenomenon Report, 2014 2H.*
[7] Wikipedia. 2016. *List of most viewed YouTube videos*https://en.wikipedia.org/wiki/List_of_most_viewed_YouTube_videos
[8] The Space. 2016. *About. http://www.thespace.org/about*
[9] Global e-Sustainability Initiative. 2015. *SMARTer2020 report.*
[10] Greenpeace. 2010. *Make IT Green: Cloud computing and its contribution to climate change.*
[11] Koomey, J. 2011. *Growth in Data Center Electricity Use 2005 to 2010*
[12]Koomey, J. Matthews, H., Williams, E. 2013. *Smart Everything: Will Intelligent Systems Reduce Resource Use?*
[13] Statistica. 2015. "*Global smartphones shipments forecast from 2010 to 2019."*http://www.statista.com/statistics/263441/global-smartphone-shipments-forecast/
[14] Statistica. 2015. "*Forecast for global shipments of tablets, laptops, and desktop PCs from 2010 to 2019."* http://www.statista.com/statistics/272595/global-shipments-forecast-for-tablets-laptops-and-desktop-pcs/
[15] International Data Corporation (IDC). https://www.idc.com/
[16] Global e-Sustainability Initiative. 2015. *SMARTer2020 report. Abatement potential calculations.*
[17] Julie's Bicycle. 2009. *Impact and Opportunities: Reducing the Carbon Emissions of CD Packaging.*
[18] Weber, C *et al*. 2010. *The Energy and Climate Change Impacts of Different Music Delivery Methods.*

[19] CIBSE. 2012. *Data centres: an introduction to concepts and design.*