

CARDIFF UNIVERSITY - CASESTUDY



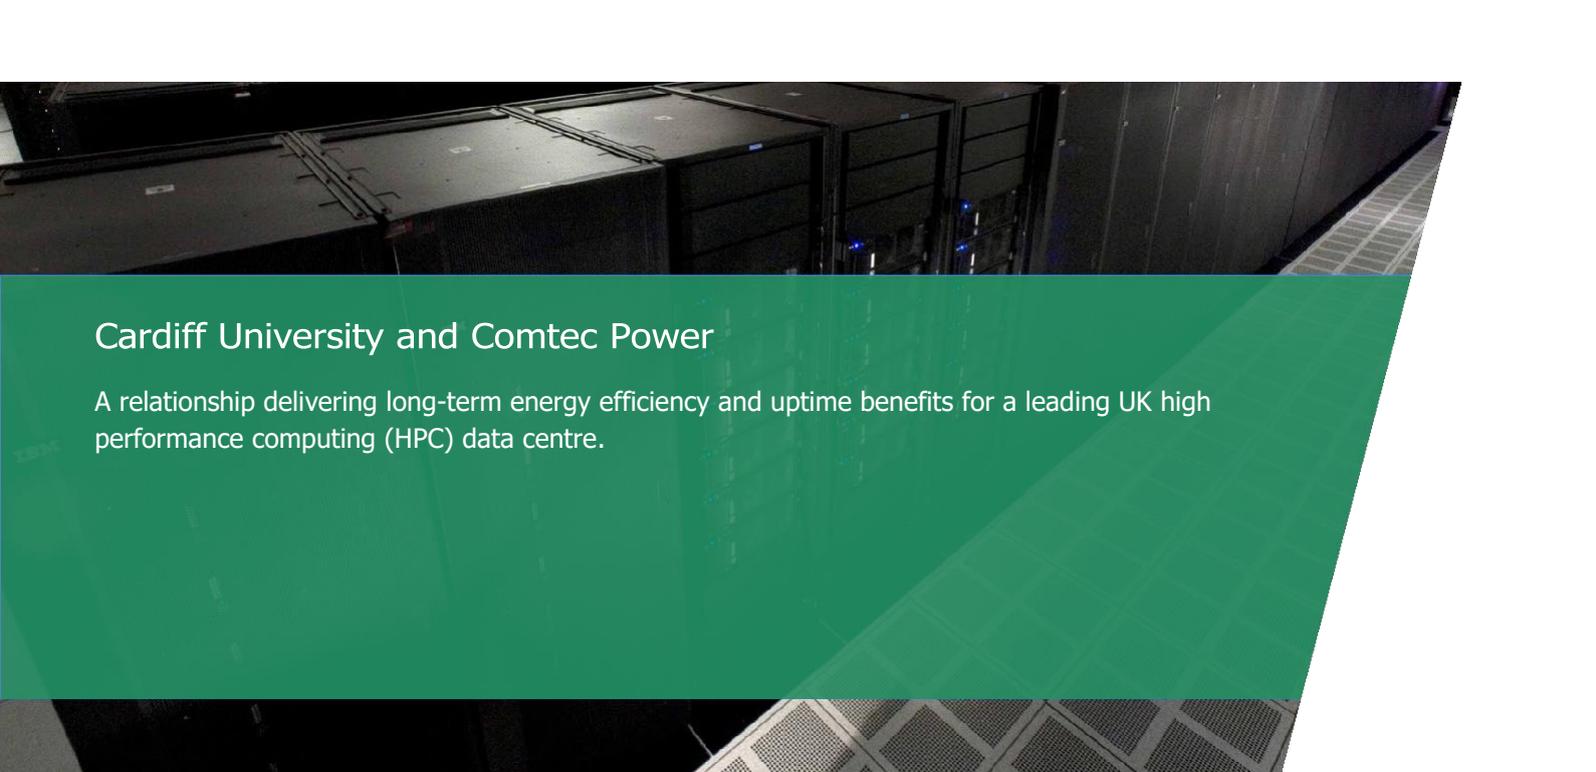
SUMMARY

Industry sector: Higher Education

Business Solution: Datacentre

Technology: APC by Schneider





Cardiff University and Comtec Power

A relationship delivering long-term energy efficiency and uptime benefits for a leading UK high performance computing (HPC) data centre.

A robust high performance computing environment supporting research excellence

Founded in 1883, Cardiff University is independently recognised as one of the UK's leading teaching and research universities, coming 5th in the latest UK Research Excellence Framework (REF). Its breadth of expertise encompasses the humanities, sciences, engineering, medicine and technology.

From the outset, the University has worked with Comtec Power, an Elite Partner to Schneider Electric, to design, construct and maintain its high performance computing (HPC) data centre.

Cardiff University's HPC data centre fulfils a number of disparate roles, from housing the servers that provide application and storage for the University's general enterprise IT needs through to hosting a range of supercomputing systems, including the latest HPC cluster, called Hawk. The cluster is operated by the Advanced Research Computing at Cardiff (ARCCA) division on behalf of both the University and Supercomputing Wales, a service which supports computationally intensive research projects across several collaborating Welsh universities – Aberystwyth, Bangor, Cardiff and Swansea.

ARCCA was originally established in 2007 to help Cardiff University maintain and build upon its position as a global centre for research.

University IT's CTO at that time, Dr. Hugh Beedie, was among the foremost advocates of green computing during the emergence of this discipline on the global stage, and frequent commentator on the subject at industry forums. He, along with the current Assistant Director of Research Computing Services, Dr. Christine Kitchen, and others, were part of the team responsible for ensuring that performance specifications for the HPC infrastructure included being as energy efficient as possible, while also being functionally advanced from a computing standpoint.

The diverse computing needs of the services operating in the Cardiff data centre imposes challenges on its support infrastructure including both power and cooling requirements, due to vastly differing operating densities of general IT and networking infrastructure compared to the supercomputing systems.

World-class technology backed by technical design and implementation expertise

Over time, Comtec has worked closely with the ARCCA team to facilitate growth and changing demands by harnessing the resilience, modularity and scalability of Schneider Electric's InfraStruxure (ISX) data centre physical infrastructure solution in conjunction with high efficiency chillers. To maximise the efficiency of the system, the strategy has been to coordinate rack density and layout with a 'close-coupled' chilled-water cooling solution and hot aisle

containment (HACS).

In conjunction with an upgrade from its original 3 x 120kw chillers to a larger array of 3 x 300kw chillers, ARCCA in collaboration with University IT and Estates department, took the opportunity to gain specific insights into the PUE (Power Usage Effectiveness) of the environment, so that the new chillers did not deteriorate energy efficiency. A critical element in managing the efficiency of the Cardiff data centre is Schneider Electric's StruxureWare for Data Centers software suite (DCIM). In use, DCIM gives insight into the power use by the data centre and the cooling capacity utilisation, allowing management to respond to changes and also to calculate metrics such as PUE in real time. This metric is a ratio of the total power consumed by the data centre to that consumed by the IT equipment alone. The closer a PUE ratio is to 1.0 the better, from an efficiency point of view.

As part of this chiller replacement programme, conducted during 2013/14, the business case worked on an assumption of an annual PUE of 1.7, which was very much a best-case scenario due to the degradation of the original chiller units. Hugh calculated that reducing the PUE to 1.4 would see the cost of the cooling upgrade comfortably pay for itself over the working 5-year lifetime of any new equipment.

As part of the same upgrade, Comtec deployed Schneider Electric's StruxureWare Data Center Operation: Energy Efficiency module as an additional component to the previously installed StruxureWare for Data Centers. Working with data inputs from extensive instrumentation that Comtec had previously installed, the new software module provided a much more comprehensive picture of power and cooling consumption throughout the data centre infrastructure and presented it on a centralised console where it could be easily viewed and analysed by the ARCCA team.

The cooling equipment upgrade together with the new monitoring software saw major improvements in the energy efficiency of the data centre, despite additional HPC servers. Depending on ambient heating conditions, the real time PUE rating has been

as low as 1.12, but more usually sits at around 1.17 – 1.20.

2017 and 2018 changes push new frontiers in ARCCA's Research Computing capability

One of the most recent evolutions in the Cardiff HPC data centre has seen the commissioning of a new supercomputer system, the third generation since ARCCA was established.

Dubbed 'Hawk', the new 8,040 core HPC cluster represents a major enhancement in resource, combining the HPC Wales¹ and ARCCA 'Raven' services into a new single supercomputing service.

Hawk consists of 7,000 cores of Intel Skylake Gold 6148 processors (2.4GHz/4.8GB memory per core, with 20 cores per processor) as the main parallel MPI partition (including a High Memory, SMP section), with an additional 1,040 cores of Intel Skylake Gold 6148 as a serial/high throughput subsystem. Accelerator technology is available through additional nVIDIA P100 GPU nodes.

Hawk is configured with 46+TB of total memory across the entire cluster, with over a petabyte of storage, allocated as a 692 TB global parallel file storage managed by the Lustre file system and 420 TB NFS/home partition for longer-term data store. Nodes are connected with InfiniBand EDR technology (100Gbps/1.0µsec latency).

The most significant impact of Hawk's introduction is its power draw, according to ARCCA Director, Prof. Martyn Guest: "The full processing load requires higher dense electrical power provisioning than the original HPC Wales infrastructure, averaging out at around 25 – 30kw per rack. This meant a number of Power Distribution Units (PDU) and Power Distribution Modules (PDM) had to be upgraded to enable the three-phase 32AMP power required. The modular/hot swappable design meant this work could be conducted without having to schedule any outages or break in service for the research community.

Prior to awarding the replacement supercomputing

¹ HPC Wales, the pan-Wales structurally funded project and predecessor to Supercomputing Wales that ran between 2010 – 2015.

contract, a feasibility study was conducted with Comtec, since some of the potential specialist HPC infrastructure from a number of suppliers had power requirements of ca. 100kW per rack. Comtec was very proactive in looking at partitioning options as this density of power would have significant implications with our UPS, power distribution and cooling infrastructure. Although this density of system was not ultimately required, it was reassuring to know some of these systems could have been accommodated, albeit with some significant investment required to modify the underpinning data centre infrastructure."

To ensure continual resilience within the environment, Comtec installed a fourth Symmetra PX160 UPS to complement the existing cluster of three systems. "The modular nature of the UPS technology means it was a simple and cost-effective way of ensuring ongoing resilience in the event of power failure, and that we didn't have to consider building a new solution to reach that scale," added Christine. "More to the point, we can keep adding to the solution to address further demands, if we need to. It is more likely that we will now reach the capacity of power to the facility rather than exceed floor space or cooling issues for the services hosted in the current data centre."

Further work was required to upgrade PDUs serving each cabinet, to ensure that sufficient power could reach the high-consumption cores. "A key consideration was the need to maintain optimum airflow behind the racks, therefore locating both PDU strips down the same side of a cabinet, and we once again found that the modular nature of StruxureWare supported this. Complementing the PDU power strips are new power breakers to cope with the increased load, and components to support the need for three-phase power in some areas, where previously we had a standard requirement for single-phase."

Comtec has also helped ARCCA develop 'zonal cooling' so that the highest densities of cooling are directed into the areas that require them. The effect of this is to mitigate unnecessary power consumption, thereby further optimising PUE. Once all the replacement services are into full production, a revised tuning exercise will be conducted by ARCCA in collaboration with Comtec, CoolTherm, University IT and Estates departments to ensure the optimum PUE efficiencies are being achieved.

Modular design enables smooth, safe transformation

Since beginning its journey with Comtec and Schneider Electric in 2006, ARCCA has gone through a steady evolution of new HPC technology to remain a UK leader in research-based computing. Throughout this period of continual change, the modular and scalable design of the data centre environment has successfully enabled success, says Christine: "These latest changes really demonstrate the beauty of taking a modular, scalable approach. We've been able to accommodate our new HPC clusters while addressing the new power load, adding UPS capacity and increasing cooling density. The only environmental constraint we anticipate is physically getting more power into our facility in the years to come. But we're confident we are being as efficient as possible in how we use the power available."

ARCCA is continually developing its long-term forward view of potential technology evolution, potentially involving cloud-bursting and colocation approaches that may even – in time – reduce the amount of power required at the current ARCCA site in central Cardiff. "Scalability and modularity works down as well as up, so we're happy we have all our options covered, whatever the future brings."

So that ARCCA can tightly monitor all the elements of its infrastructure to ensure maximum efficiency, new software – in the form of Schneider Electric's EcoStructure IT Expert – is also being deployed to enable a complete view of the entire multi-vendor data centre environment.

Calculating PUE accurately for a data centre with such mixed functions as Cardiff's presents its own challenges. For example, cooling provision for the systems supporting general IT needs remains reasonably consistent, whereas for the HPC work that ARCCA performs the systems, when they are running, tend to be operated at peak power.

"By the end of 2018, we will have the new software in place to complete a PUE assessment with all the new Hawk infrastructure. As well as kicking off the real-time monitoring process with EcoStructure IT Expert's enhanced capabilities, it will also give us a 'cold weather' wintertime benchmark to compare against hot weather readings we plan to collect next summer."

A deep relationship built on trust and competence

As well as providing expert advice at every step of the journey, Comtec delivers a programme of preventative maintenance to underpin ARCCA's commitment to resilience and uptime. Regular onsite checks of electrical systems, such as UPS battery testing and replacement, are complemented by periodic software upgrades to the latest version of Schneider Electric's DCIM technology. Comtec also provides a comprehensive service level agreement for the maintenance of equipment components, including next business day hardware replacement and remote monitoring of the ARCCA environment with skilled engineers.

"We invest a lot in being the supercomputer experts

so that we can address all sorts of challenges from research projects that continually push the boundaries of what's possible. That's what makes Comtec such a valuable partner. We rely on their expertise to back us up; guiding our decision making on anything electrical power and cooling related.

The result of this long-term relationship with Comtec is our ability to evolve and transform our own capabilities in the confidence that our data centre environment will evolve and transform with us. And throughout all this time, with Comtec at our side, we've kept our services running, achieving outstanding uptime levels and keeping pace with our customers' expectations for 24/7 operations."

ABOUT COMTEC POWER

Comtec Power specialises in the design, build, maintenance and on-going management of datacentres and server rooms.

Our service encompasses new builds as well as the refurbishment and upgrade of live facilities. Our goal is to create efficient, technically excellent facilities that are scalable to meet the power and cooling demands of latest generation IT equipment.

We are an experienced systems integrator with a thorough knowledge and understanding of the IT industry. We are perfectly positioned to deliver highly resilient solutions for high-density IT deployments and our on-going maintenance and management services help to reduce the threat of systems downtime, mitigate the risk of thermal shutdown and keep facilities operating optimally.

Head Office
Comtec House
Albert Road North
Reigate
Surrey RH2 9EL

T: 0800 008 7599
E: info@comtec.com
www.comtec.com

