

Free cooling on the Mediterranean shore: Energy efficiency upgrades at PIC

V Acín^{1,4}, M Delfino^{2,3,4}, A Herbera² and J. Hernández²

¹Institut de Física d'Altes Energies (IFAE), Edifici Cn, Universitat Autònoma de Barcelona, E-08193 Bellaterra (Barcelona), Spain

²Universitat Autònoma de Barcelona, E-08193 Bellaterra (Barcelona), Spain

E-mail: delfino@pic.es

Abstract. Energy consumption is an increasing concern for data centres. This paper summarizes recent energy efficiency upgrades at the Port d'Informació Científica (PIC) in Barcelona, Spain which have considerably lowered energy consumption. The upgrades were particularly challenging, as they involved modifying the already existing machine room, which is shared by PIC with the general IT services of the Universitat Autònoma de Barcelona (UAB), with all the services in full operation, as well as the introduction of “free cooling” techniques in a location 20 km from the Mediterranean Sea. The upgrades targeted three distinct areas: First, the segregation of hot and cold air zones using an innovative horizontal layout, where hot air is channelled through openings in a false ceiling to a second story hot air plenum. This segregation allows increasing the cold air inlet temperature according to the latest ASHRAE recommendations. Second, the introduction of an outside air economizer which replaces obsolete CRAH systems with air-to-air heat exchangers. This system, built entirely from industrial components, also incorporates an adiabatic cooling module and enables the “free” removal of over 300 kW of IT heat load during 6000 hours a year. Third, the introduction of UPS systems based on IGBT technology, in order to better match the impedance characteristics of the IT load. In addition, a transversal activity has been done to fully integrate cooling and UPS infrastructure monitoring into PIC's overall IT monitoring framework based on Nagios. This required development of a ModBus/TCPIP gateway server.

1. Introduction

Energy consumption is an increasing concern for data centers. There are several aspects to these concerns:

- Total energy consumption and its impact on overall costs
- The fraction of energy consumed which ends up powering computing equipment. This is usually measured using the PUE parameter [1], defined as the ratio of total energy consumption to computing equipment energy consumption.
- High PUE values imply high energy consumption by the cooling system. If the cooling system reaches its absolute capacity, it is impossible to grow the consumption of the computing equipment.
- Ecological concerns, such as the carbon footprint.

³ Author to whom any correspondence should be addressed.

⁴ Also at Port d'Informació Científica (PIC), Campus UAB – Edifici D, E-08193 Bellaterra (Barcelona), Spain



The Port d'Informació Científica (PIC) delivers Spain's contribution at the Tier-1 level to the Worldwide LHC Computing Grid [2] and serves as main data repository (Tier-0) to a number of projects in astrophysics and observational cosmology.

PIC's computer room installation is very compact, consisting of two distinct computer rooms.

1.1. The main computer room

The main machine room consists of a 140 m² partition within a 400 m² traditional raised floor computer room built in the late 1980s. The rest of the computer room is used for general services of UAB. Until the upgrades described in this paper, the computer room had a high ceiling with no separation of hot and cold air, and the cooling was based exclusively on chilled water generation.

The impact of PIC's initial 2003 installation in the building was substantial. Installed power for computing equipment was increased from 40 to 200 KVA. Although the PUE was not directly measured at that time, engineering studies at that time showed that only minor modifications needed to be performed on the cooling system, as it had a large cold water generation capacity originally destined for cooling an IBM 3090 class mainframe.

Given the ever growing power needs of the installation, a major upgrade of the power input and diesel backup was undertaken in 2008 by the UAB Architecture and Logistics Department. Two new transformers of 1 MVA each were installed as the main power source, backed up by up to four 500 KVA diesel generators. Unfortunately, just as the power upgrade was being finished, a first measurement of the PUE yielded a value of 2.3. This very high value meant that total power consumption would become unaffordable, and that the cooling system was at the limit of its capacity.

Part of the high PUE value was due to the building architect's choice of placing the chillers in the basement, requiring about 100 kW of power to fans to bring outside air to the basement. In addition, the chillers were 25 years old and had relatively low energy efficiency. A program was started in 2010, co-funded by UAB, IFAE and European Union ERDF funds, to install two new chillers on the roof of the building. The result was a lower PUE of 1.7 by the beginning of 2012.

Nevertheless, a PUE of 1.7 is still high and generates substantial costs, especially considering Spain's high electricity prices. Hence, the additional program described in this paper to lower the PUE and allow increased capacity was executed in 2012-2014, funded by the same sources.

1.2. The compact data centre module

A separate, higher efficiency, 25 m² computer room module [3] was quickly built in 2010 as a stop-gap measure to handle 80 KVA of additional load needed by the Spanish WLCG Tier-1 center. Built out of standard industrial components, this module implemented an almost complete separation of hot and cold air. Measurements show that the PUE of this module is around 1.6 to 1.7, and that it is almost exclusively determined by the characteristics of the air conditioning system, which is based on gas expansion. The module has worked well for four years, and allowed to lower the pressure on infrastructure in the main computer room in order to ease the installation of the upgrades described in this paper.

2. Computer room refurbishment

2.1. Vertically stacked hot air containment scheme

Best practices in modern data centers and the experience gained with the compact data centre module made it clear that separation of hot and cold air should have maximum priority. The challenge was how to implement it, as the modifications had to be done with the PIC and UAB data centres in full operation, with a maximum of 24 hours of scheduled downtime.

Hot and cold air containment schemes were studied. Two strong constraints were found:

- PIC's tape robots share the same space as the rest of the computing equipment and their temperature should not exceed 32°.

- There is a single common set point for the cold air temperature in PIC's and UAB's computer room, due to the fact that cooling equipment is fully shared by both machine rooms.

In addition, PIC and UAB management desired to raise the cold air temperature as much as allowed by the updated ASHRAE standards [4], in order to lower cooling expenses further.

Measurements were made by PIC and UAB staff and compared to Computational Fluid Dynamics (CFD) simulations to better understand air flow as a function of temperature. The engineering firm PQC [5] was hired to explore alternatives and provide a baseline design (including the introduction of free cooling, see the next section) that could be tendered. The conclusions were:

- Cold air containment was not suitable unless hot air containment was also implemented, increasing the cost of the upgrade.
- Guiding hot air to the cooling units placed at the perimeter of the machine room was crucially important.
- Reversing the air flow, in order to use the false flooring to contain the hot air, was complex and risky, and could end up with limitations on how much hot air could be handled.
- A vertically stacked hot-air containment option was identified as the best option.

This vertically stacked containment scheme is illustrated in Figure 1. A false ceiling spanning the 400 m² of the entire computer room was built, creating a volume where hot air is trapped and guided towards the cooling units. The false ceiling is made of modular panels which can be solid or grills. Grill panels are placed above the hot air aisles, allowing air to rise into the containment volume.




Figure 1: Vertically stacked hot air containment scheme

The gap between the top of the racks and the false ceiling, as well as the vertical space at the edge of rack rows, is closed by plastic lame curtains hanging from the ceiling in order to delimit the hot air aisles. Photographs of the physical layout around the racks before and after the upgrade are shown in Figure 2. Most of the computer room volume is filled with the uncontained cold air, providing a large plenum from which equipment can easily draw air with minimum fan speed, helping to save energy.

2.2. Increase of cooling capacity through introduction of free-cooling

As was mentioned in the introduction, most of the cooling equipment in the original computer room was obsolete and there was a need for increased cooling capacity. The main Computer Room Air Handlers (CRAHs) were replaced with new units which are integrated with a new indirect free-cooling system and with the existing chillers.

Indirect free-cooling (more properly known as outside air economizer) is a technique by which computer room air is circulated through an air-to-air heat exchanger, transferring heat to exterior air without mixing. The air-to-air heat exchanger is based on a set of metal plates. Each side of a metal plate is exposed to one of the air circuits, and heat is transferred by conduction within the plate. The term indirect is meant to emphasize that the computer room and exterior air do not mix, which avoids problems of humidity and air-quality control. This is especially important in PIC's environment,



Figure 2: Original open high-ceiling room (left) and detail of false ceiling with grill over hot aisle delimited by plastic lame curtains (right)

where high humidity conditions occur often and the surrounding area has greenery which can produce large amounts of pollen and loose leaves at certain times of the year.

Free-cooling is usually associated with installations at high latitudes with rather cool weather. Thus, PIC's choice of installing a free-cooling system at a location with 41° latitude and near the Mediterranean Sea may seem surprising at first. The 41.5°N 2.1°E location, however, is West-Northwest of the city of Barcelona, separated from the sea by the Collserola range and only 90 km from the foothills of the Pyrenees mountain range. This creates a micro-climate with substantial temperature drops at night throughout the year, and cooler temperatures than Barcelona in fall, winter and spring. Simulations using historical weather data show that weather conditions will allow some amount of free-cooling up to 6000 hours per year.

One key design feature is to allow dynamic cooperation between the free-cooling scheme and the chilled water backup system. The bid to install the system was awarded to Agefred [6] which proposed using equipment by Tecnivel [7]. A computer generated drawing of a free-cooling unit is shown in Figure 3. The diamond-shaped object at the centre mechanically regulates the area of heat exchanger plates exposed to the air flow.

The free-cooling stage is used to cool the computer room air as much as the exterior temperature will allow, and then the air is passed through a panel with a regulated chilled water serpentine to bring it to the set point. In addition, when the exterior air is sufficiently hot and dry, an adiabatic unit is used to spray water drops whose evaporation can lower its temperature by as much as 5°C, thus extending the range over which free-cooling contributes.

One issue is that the free-cooling units are rather bulky. The three free-cooling units needed for the PIC and UAB upgrade were installed in reclaimed office space located at one of the sides of the computer room. Installation required temporarily removing an outside wall in order to bring the equipment into the building, as shown in Figure 4. The exterior air interface of the free-cooling units was done by replacing three existing windows with air grills for the intake and by building three white cubes for the exhaust. The aesthetics of the building were thus preserved.

The three free-cooling units draw computer room air from the hot air containment space above the new false ceiling (see section 2.1), cool it to the set point, and inject it into the pre-existing false floor, providing the functionality indicated by the striped arrow in Figure 1.



Figure 3: Free-cooling equipment from Tecnivel. Exterior air flows through the top and computer room air through the bottom.



Figure 4: Photograph of the arrival of one of the free-cooling units. The reclaimed office space can be seen through the opening left by the removal of the outside wall. The computer room occupies the area directly on the left, with a sloping roof rising towards the left.



Figure 5: Photograph of the facade after the installation was completed. Intake air grills replaced three windows and white cubes were built for the exhaust vents.

No relocation of racks or equipment was needed, since the underfloor cold air scheme was not changed. The ceiling was constructed and the free-cooling units were pre-installed in the reclaimed office space with both computer rooms fully functional. The yearly 24 hour scheduled downtime was used to remove the old CRAHs, couple the free-cooling unit air ducts and interface the chilled water. The plastic lame curtains were installed progressively after the shutdown, with both computer rooms fully functional again. Thus, this major cooling system upgrade was accomplished with no additional downtime.

2.3. Monitoring system integration

For the last five years, PIC has made a major investment in integrating power consumption monitoring into its standard Nagios/Ganglia monitoring setup. The long term strategy is to build energy awareness into PIC's highly automated operations environment.

The Tecnivel free-cooling equipment is built out of industrial components and sensors. Monitoring information is available via industry-standard ModBus. An interface has been developed to be able to read from ModBus via Ethernet and TCP/IP and feed the information into PIC's Nagios/Ganglia setup. An example showing temperatures from various parts of the system is shown in Figure 6.

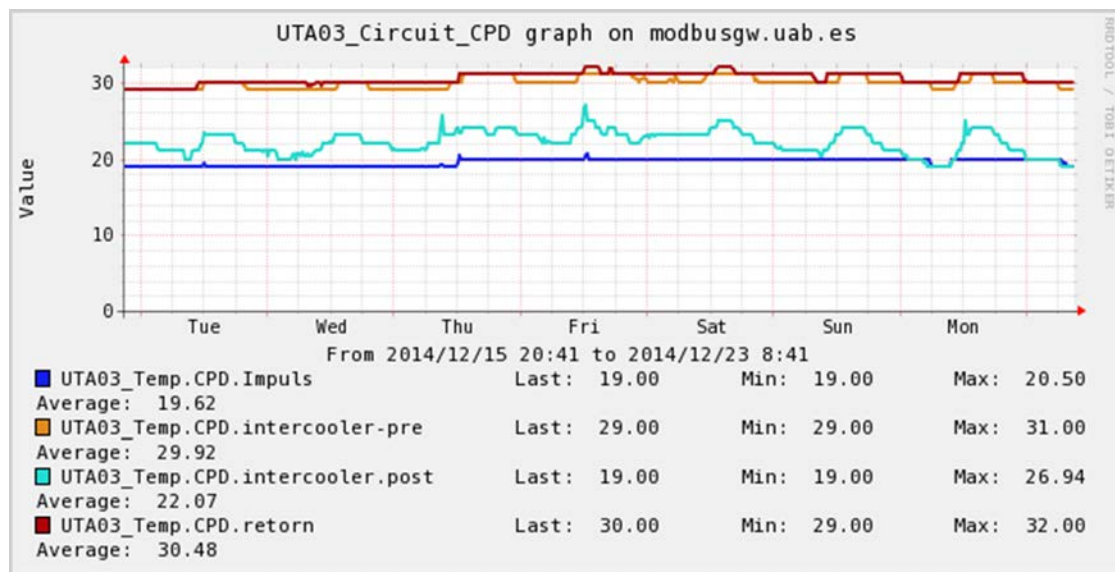


Figure 6: Ganglia plot showing temperature readings in mid-December 2014. The measurement points are: Intake from computer room (red/orange), exhaust from free-cooling stage (turquoise), final cold air delivery to computer room (dark blue).

3. Preliminary results on energy efficiency improvement

The installation plan for the upgrade had its maximum priority on minimizing downtime and completing the hot-air containment installation. Fine tuning of the free-cooling system will be done over a one year period, supported by analysis of the sensor information. Thus, only preliminary results are reported here, but they clearly point to a high degree of success of the upgrade.

The following qualitative or semi-quantitative observations have been made:

- Thanks to the hot air containment, temperature variations in the non-contained cold air volume have notably diminished
- Thanks to the plentiful supply of uncontained cold air, internal fan speeds in power-hungry computing equipment have dropped to the minimum setting

Progress is being made towards a quantitative understanding of the behaviour of the system. The information in Figure 6 can be used to visualize the savings brought by the free-cooling system. The overall cooling needs are proportional to the area between the red/orange and dark blue curves, corresponding to the computer room hot and cold air set points, respectively. The area between the red/orange and turquoise curves represents the contribution of the free-cooling system, with the rest of the heat being absorbed into the chilled water system. As can be seen towards the right side of Figure 6, there were several hours on the night from Sunday to Monday where the free-cooling system was able to handle the cooling load completely.

There are, of course, many details to understand, in order to accurately measure the new PUE. The free-cooling system does consume some power (mostly in fans), etc. Nevertheless, a first measurement of the winter-time PUE has been made yielding a value around 1.3. An extrapolation using weather data indicates that the yearly average PUE will be in the 1.4 to 1.5 range. This improvement of 0.2 to 0.3 units in PUE translates into energy expenditure savings such that the roughly 400 thousand euro investment will be recovered in four years. In addition, the requirement of having additional cooling capacity in order to allow expansion of the computer centres has been fulfilled.

Work will continue during 2015 to fine tune the system. An additional source of savings comes from the raising of the set point for underfloor cold air delivery from 13 to 21°C, substantially reducing the load on the chillers. This is made possible by the large plenum of cold air for the equipment which has been mentioned above.

4. Uninterruptible Power Supply with improved energy efficiency

PIC's Uninterruptible Power Supplies (UPS) setup consisted of a 200 KVA UPS for the main computer room and an 80 KVA UPS for the data centre module. The upgrade program will replace these with a 500 KVA modular UPS based on Insulated Gate Bipolar Transistors (IGBT) [8]. UPS equipment has higher efficiency when used near its full capacity. PIC's new 500 KVA supply will have two 250 KVA modules, allowing to keep one module near full capacity with a base load and using the other module to provide power for peak loads. In addition, the IGBT technology provides higher conversion efficiency than the equipment it replaces. This part of the upgrade was being installed at the time of the CHEP 2015 conference.

5. Conclusions

An ambitious upgrade of the computer room shared by PIC and UAB has been successfully undertaken. An innovative vertically stacked hot air containment scheme has been chosen, which provides a number of benefits. An indirect free-cooling system, based on industrial components, has been deployed using a cooperative scheme with the standard chilled water system. Early measurements indicate a winter PUE of 1.3, leading to estimates of a yearly average PUE of 1.4 to 1.5. This represents a substantial increase in energy efficiency, generating savings which will recover the investment in four years, and will allow an increased capacity of the computer centres. Installation of the air containment and free-cooling system was done with the computer centres in full operation and profiting from a yearly 24 hour scheduled downtime, without incurring any additional interruption.

Acknowledgments

The Port d'Informació Científica (PIC) is maintained through an agreement between the Institut de Física d'Altes Energies (IFAE) and the Centro de Investigaciones Medioambientales, Energéticas y Tecnológicas (CIEMAT), with additional support from the Universitat Autònoma de Barcelona. IFAE is funded by the Generalitat de Catalunya as part of the CERCA research centres. CIEMAT is funded by Spain's Ministerio de Economía y Competitividad.

The support of the ERDF fund is gratefully acknowledged.



References

- [1] The Power Usage Effectiveness (PUE) was first defined by the Green Grid organization. See <http://www.thegreengrid.org/sitecore/content/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE.aspx> (accessed on 17/05/2015)
- [2] Lamanna, M 2004 The LHC computing grid project at CERN *Nucl. Inst. and Meth. in Physics Research* **534** 1
- [3] Acin V, Cruz R, Delfino M, Martinez M, Rodriguez M and Tallada P 2011 First operational experience from a compact, highly energy efficient Data Center Module *J.Phys.Conf.Ser.* **331** 022001 DOI: 10.1088/1742-6596/331/2/022001
- [4] ASHRAE Technical Committee 9.9 2012 Thermal Guidelines for Data Processing Environments, Third Edition (Ashrae Datacom) *American Society of Heating, Refrigerating and Air-Conditioning Engineers* ISBN-13: 9781936504336 ISBN-10: 1936504332
- [5] PQC, S.L. <http://www.pqc.es>
- [6] Agefred, S.L. <http://www.agefred.es>
- [7] Tecnivel, S.L. <http://www.tecnivel.es>
- [8] Baliga, B J 2015 The IGBT Device: Physics, Design and Applications of the Insulated Gate Bipolar Transistor *Elsevier* ISBN: 978-1-4557-3143-5