

ENERGY EFFICIENCY IN DATA PROCESSING CENTERS

Redondo Gil, Carlos ^{(1) (2)}

⁽¹⁾ Scientific Manager, **Castile and León Technological Center for Supercomputing (FCSC)**
Edificio CRAI-TIC, Campus de Vegazana s/n.
University of León. E-24071-León (Spain).

carlos.redondo@fcsc.es <http://www.fcsc.es>

⁽²⁾ Electrical Engineering & Systems Engineering and Automatic Control Department
Faculty of Industrial and Computer Engineering, **University of León**. E-24071-León (Spain).

carlos.redondo.gil@unileon.es <http://www.unileon.es>

Abstract

On a global scale, IT equipment is responsible for 2% of CO₂ emissions, the same quantity as that emitted by the totality of airline companies world-wide.

Furthermore, and as a point of reference, electricity consumption by data processing centers in the USA constitutes approximately 2% of total national electricity consumption. This high rate of consumption implies that IT is currently the second most contaminating industry on the planet, surpassed only, once again, by the aeronautical industry. However, given its present rate of growth, IT will shortly head the list.

This article assesses the present and future situation of IT center energy consumption and associated environmental effects, and also looks at how state-of-the-art technology, correctly implemented, could ensure significant rationalization of data processing center energy consumption. The article will examine optimization techniques, specific problems and case studies.

KEY WORDS


Energy Management. Data Processing Centers. Energy Efficiency.

1. Introduction

1.1. A global problem

Since the inception of Computing, one of the principle driving forces behind the growth of this branch of engineering has been the development and assembly of increasingly powerful systems, capable of resolving ever more complex scientific problems and processing an ever greater volume of data in less and less time. One consequence of this has been the proliferation of the large data processing centers (DPC) associated with research institutions (BSC, FCSC, CeSViMa, CESGA, CIEMAT, CiCA,...) and companies working in telecommunications, banking and energy (Telefónica, BBVA, Repsol, etc.).

The pie chart below gives an estimate of DPC energy consumption, by component, in the USA. It is predicted that in 2011, this kind of technological installation will represent 10% of electricity consumption.




Source: Gartner Group 2007

Figure 1. CIT as a causative factor in Global Warming.
Percentage of energy consumption each ICT

Within a DPC there are two principle sources of energy consumption. One comprises the information systems themselves, including the hardware, data storage systems and network interfaces. The other source comprises the cooling systems which are essential to guarantee the proper performance of information systems. Clearly, cooling system consumption is closely related to information system consumption, as the more the latter are used, the more heat they generate, thus increasing cooling requirements. In the case of either system, the direct result of energy savings would be a reduction in both running costs and environmental impact [2,3,5,11,12].

DPC consumption world-wide is over 61,000 million Kw*hr per year. However, if energy efficient solutions were deployed, not only would annual savings of over 4,000 million euros be achieved, but, in the USA alone, the reduction in CO₂ emissions would be equivalent to taking five million cars off the road. Therefore, it is becoming increasingly imperative to implement solutions which optimize DPC energy consumption. In this vein, it is noteworthy that over the last ten years this consumption has multiplied ten-fold, with cooling systems being responsible for over 60% of this consumption. It is this aspect which is especially

susceptible to implementation of efficient models of technological innovation.



Source: EPA Report to Congress on Server and Data Center Energy Efficiency. August 2, 2007.

Figure 2. DPC energy consumption predictions for the USA, showing various scenarios.

A whole series of remedial innovative initiatives are emerging in an attempt to solve this problem:

Cisco, Microsoft and the Reykjavik electricity company are investigating ways to exploit electricity generated by the volcanic country's geothermal springs, in order to supply large server farms¹.

Google may be considering relocating its servers out at sea, on ships or platforms anchored 12 kilometers from the coast. The idea is to reduce energy consumption by using sea water to cool the data center².

In 2003, a server cost approximately 30,000 euros. These days, a machine with a much higher processing capacity can be bought for 5,000 euros. Thus, the cost of electricity and the cooling system required to maintain optimal server performance has begun to outstrip the price of the server.




Figure 3. Server parks (in millions of servers)

¹ www.elpais.com/articulo/internet/Islandia/quiere/ser/servidor/elpeupes
² http://technology.timesonline.co.uk/tol/news/tech_and_web/the_web/article4753389.ece




Figure 3. Cost of server parks globally

The price of servers is becoming equal to that of their energy consumption. In 2009, 48% of the technology budget was allocated to paying for electricity. Electricity and cooling costs have increased eight-fold, whilst the price of a server is a tenth of what it was in 1998.

59% of the energy consumed in a data processing center is accounted for by the running of the machines. Studies have established that throughout the world, the remaining 41% is used as follows:

- Between 23 and 54% on cooling systems.
- Between 6 and 13% on complementary systems (generators, uninterruptible power supplies (UPS)).
- Between 1 and 2% on lighting.
- Between 38 and 63% on infrastructures.




Figura 4. Power flow in a typical Data Center

At present, because the government subsidizes energy costs, the cost of a Kwh is relatively low, approximately 12 cents/Kw*h. However, the day the subsidies dry up, energy consumption will become one of a DPC's most expensive outgoings.

Recently, **Fujitsu Siemens** carried out research into data centers in the UK. According to the report, these centers house some 880,000 servers, generating 1.2 million tons of CO₂ per year due to bad configuration, infra-utilization and excessive use of air conditioning. These figures are higher than the annual emissions of a city such as Oxford (981,000 tons).

1.2. The perspective in Spain.

Spain enjoys the second most generous emissions quota under the Kyoto Protocol, with permitted emission levels of 15% more than those recorded for 1990, whilst the European Union's objective is an 8% reduction on 1990 levels. The most recent data available indicates that Spanish emissions are 43% higher than 1990 levels, almost three times the limit imposed in the Kyoto Protocol.

Furthermore, Spain is behind the rest of Europe in terms of compliance with emissions commitments. According to environmental and energy regulations published by the European Community, Spain will be required to increase renewable energy production by 42%.

Of all industrialized countries, Spain is the one where greenhouse gas emissions have increased most since 1990, the year taken as a reference point in the Kyoto Protocol. Even worse is the fact that despite promises of improvement, each year is worse than the one before. In 2004, the volume of polluting gases released into the atmosphere was 47.87% higher than in 1990.

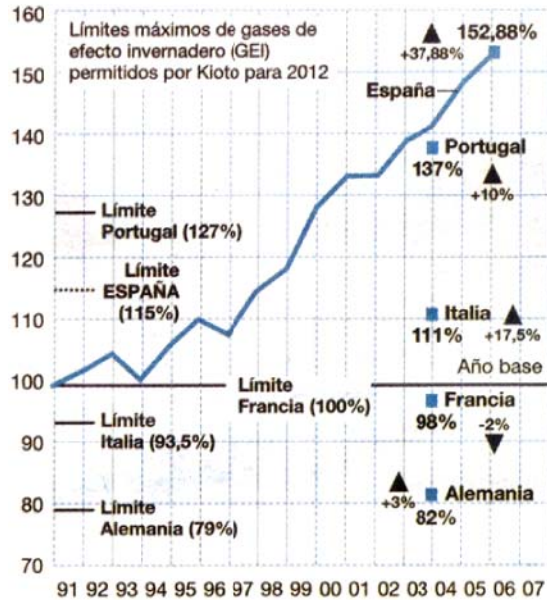


Figure 5. Evolution of Greenhouse Gas emissions

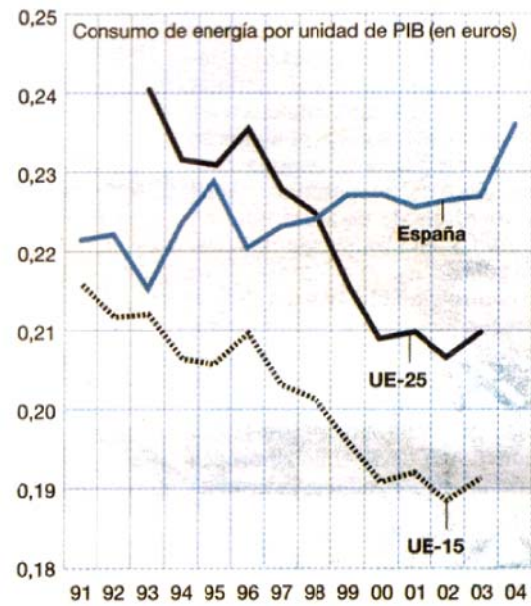


Figura 6. Non-compliance with the Kyoto Protocol: Kyoto only permits a 15% increase
Source: *El Mundo* (Spanish daily national newspaper)

2. Energy management model (case study: Caléndula – Centro Tecnológico de Supercomputación de Castilla y León – Castile and León Technological Center for Supercomputing)

One of the most active non-profit making organizations, the Green Grid, has proposed a metric which is becoming the standard when measuring this type of efficiency. It comprises two different parameters, and two equivalent forms of measurement: PUE (*Power Usage Effectiveness*) and DDCiE (*Data Center Infrastructure Efficiency*) [14].

PUE represents the total power consumed by the installations, divided by the power consumed by the IT equipment. A PUE of 2 means that 50% of DPC energy consumption corresponds to IT equipment, and the other 50% to complementary resources and cooling. The DDCiE, meanwhile, represents the quotient between IT equipment power x 100, divided by total power supplied to the DPC. The result is a percentage, where the higher the figure, the more efficient the energy consumption. Thus, for example, a DDCiE of 50% is equivalent to a PUE value of 2.

With energy efficiency in mind, the Castile and León Supercomputer Center has designed a server system which has not been adopted to date in similar facilities. Not only was the computer calculating power considered, but also an energy consumption objective which makes the Center the most energy efficient in Spain.

The Center's mission is to provide a central reference point for the application of the latest-state-of-the-art ideas and developments in these systems.

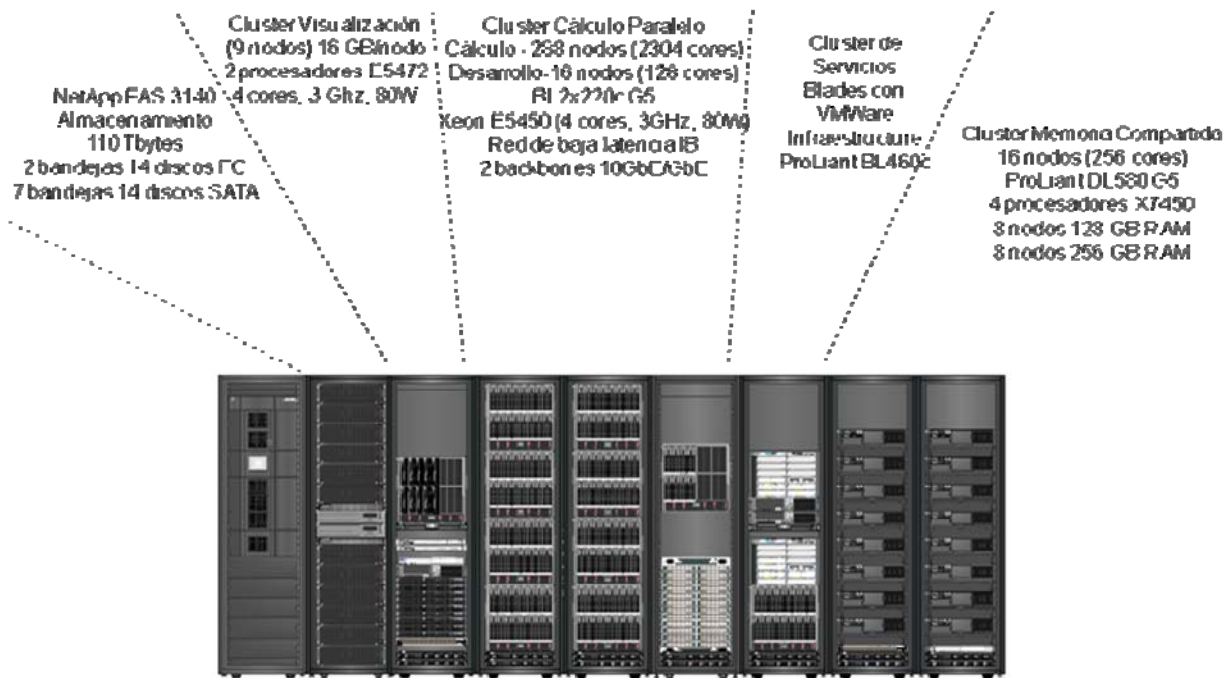


Figure 8. Cutting edge technological infrastructure for trial and assessment.
Centro de Supercomputación de Castilla y León

2.1. Description and aims of the Foundation

The Foundation is a public body created in 2008 in the Autonomous Region of Castile and León, and is integrated into the Spanish e-science Network. Its mission is to improve research facilities at the University, at Research Centers and at companies within Castile and León by:

- Fomenting innovation in the Knowledge Society:
 - in the field of intensive calculation,
 - communications, and
 - Advanced Services.
- Contributing to economic development in the Region through technological improvement, and to improving the competitiveness of companies through Innovation and Technological Development and supporting Basic and Applied Research.

2.2. Calculation Infrastructures

The heart of the Foundation is a super computation unit installed in the Center for CIT Resources for Learning and Research at the University of León. The unit contains the following features: Parallel computation cluster, Visualization cluster, Share memory system, Development nodes.

The Equipment described provides a calculating capacity of 30 Tflops, situating it in second place among all such centers in Spain³, and among the top 150 world wide⁴.

A key feature to the design of this supercomputer has been energy efficiency. Electricity consumption is the decisive factor in running costs, and represents a growing concern due to its environmental impact. The supercomputer's electrical power is approximately 120 Kw, a long way from the 683 Kw of the "Mare Nostrum" or the 267 Kw of the "Finis Terrae", placing it amongst the first 5 installations world wide which have overcome the 500 Mflops/w barrier.

The Center in León center has a 170 Kw capacity, for a calculating power of 33 Tflops.

The space occupied by the supercomputer, 90 m², also constitutes a technological challenge of the first magnitude, as the installation of such a high calculating power in such a reduced space necessarily implies the need to dissipate the great heat generated by the machines rapidly. This is only possible thanks to the installation of an efficient



Figure 9. Conceptualization of Research, Development and Innovation
Source: Own work

³ The center occupying first place is the "Mare Nostrum" in Barcelona with a capacity of 94 Tflops. The CESGA "Finis Terrae" installation in Galicia has a capacity of 21 Tflops.

⁴ Top500 June 2008, 53rd place world wide and 131st in May 2009.

water cooling system, which exploits the low external temperatures which characterize the city of León throughout most of the year.

4.5 million euros have been invested in the project to date, of which 3.3 million euros were allocated to the acquisition of computing equipment supplied by Hewlett Packard, Bull and Satec, and 1.2 million euros were spent on equipping the space where the supercomputer was to be housed, and on buying and installing the infrastructures necessary for operation (electricity, air conditioning, security, fire extinction systems, etc)

2.3. Infrastructure design criteria

To move from “room cooling systems” to “dynamic smart (rack) cooling”, to use cooling economizers, to scale appropriately through the use of modular power and cooling equipment, to employ high efficiency UPS and to implement dynamic control of the cooling plant.

2.3.1. Hot and cold aisles technique

The **hot and cold aisles technique** was implemented, which prevents the hot air leaving the machines from mixing with the cold air being generated to cool others.

The principle cause of DPC cooling inefficiency is allowing the air which has passed through one device – and therefore has a higher temperature, as it has extracted the heat from the interior of the device – to mix with the cold air from the cooling units. This leads to a reduction in cooling performance, as room temperature rises and capacity for cooling the units drops. In turn, efficiency is reduced due to a lower differential between input/output temperatures.




Figure 10. Graphic of air flow in an installation such as that at the FCSC.

In addition, fluid mechanics concepts have been applied in the implementation of configurations and designs to ensure that air flows in the most efficient manner possible according to the thermal load at any particular moment.

2.3.2. Elimination of “hot points”

The temperature and humidity requirements of data centers are conditioned by the characteristics of the most demanding units. Thus, the existence of hot points dictates cooling for the entire installation. If these “hot

points” were eliminated, producing a constant thermal load in the center, this would create a relative reduction in temperature and humidity requirements, thus contributing to a reduction in energy required for cooling the center.

2.3.3. Rack cooling. Water cooling. “In-row”

Cooling the rack on which the servers are mounted, in order to eliminate heat as close to the source as possible. Move from room cooling to dynamic rack cooling. The expression “**In-row**” is applied to new techniques being developed by CIT unit manufacturers, which generally comprise the inclusion of **cooling as an integral (internal) part of the unit**. Just as a unit incorporates a supply source, so too can a cooling system be integrated.

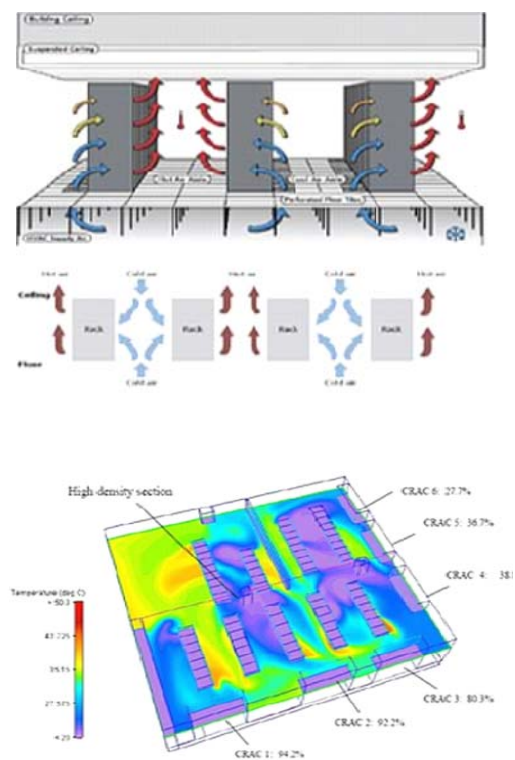


Figure 11. Diagram of temperature distribution in a room with floor cooling (traditional cooling system)

Water cooling. Instead of cooling the equipment with air, the units can be cooled directly using chilled water circuits. Use of **cooling economizers**. This cooling system provides a more efficient thermal circuit, enabling water flow to operate at temperatures higher than would be possible with the air alternative, and thus increases the possibility of using “free cooling”.

2.3.4. “Free Cooling”

“Free cooling” is possible when the external temperature is very low in comparison with that of the interior. It consists in exploiting external air to

cool the DPC, thus obviating the need to use the chilling systems

Energy savings through the use of “free cooling” depend on the climate of the DPC location, and as would seem obvious, the colder the climate, the higher the energy savings. This is one of the principle arguments currently being applied when selecting sites for new DPCs.

To scale appropriately through the use of modular power and cooling equipment.

2.3.5. Virtualization

At present, one of the best solutions for preventing a large proportion of the CO2 emissions generated by data centers – apart from improving cooling systems - is server virtualization.

Virtualization creates a simpler and more intuitive virtual reality parallel to the real world, capable of providing technological services to organizations, the management of which is more agile thanks to a “software layer” which is then transferred to the real world.

One of the consequences of virtualization is a drastic reduction in the number of units needed in a server farm, with all the associated savings in maintenance and energy costs and reductions in the emission of toxic gases that this implies.




Figura 12. VMware infrastructure. Enterprise virtualization

In this way, a company’s computing resources are used in a more efficient manner whilst at the same time, offering greater flexibility in the management of such resources.

2.3.6. Modular design

To establish independent compartments for the installation of CIT units, in such a way that the operation of one does not affect the others, so that each of the rooms can be adapted to the scale required at any specific moment, rendering DPC resources responsive to actual needs.

Similarly, it should be pointed out that any electromechanical device that generates heat should be housed away from CIT units in order to avoid overloading the cooling systems.

3. The AECePro (HPC: Energy Savings and Management in the HPC Environment – Environmental Management) Project.

For reasons of economy and improved performance, DPCs generally use a series of servers configured as a computer cluster, using a workload distribution application to plan user task execution. Frequently, demand for calculating power in these centers alternates between periods of intense activity, with high demand for use, followed by periods of low demand. Nevertheless, the scale of the computing systems usually corresponds to that necessary during periods of maximum activity. Furthermore, the size of the cooling system is also designed to ensure safe operation of the cluster during periods of maximum activity, allowing for the emergence of hot points in specific areas of the data center due to the concentration of tasks in certain cluster components.

The technological basis of the AECePro project is centered on disconnecting those computers which are unnecessary for DPC performance at any given moment. By then distributing the tasks carefully among the active computers in order to avoid hot points, it becomes possible to save energy, both in terms of what the computer system components consume directly, and in terms of the energy required by the cooling system. Although these measures are already implemented manually or semi-automatically in some centers, the aim of this project is to design an application which automatically analyzes user needs and appropriate parameters for defining DPC performance, activating only those computer system components necessary at any given moment, and distributing tasks efficiently.

The principle is simple, but instantiation in the form of an effective computer application implies finding the answer to complex questions such as computer system monitoring, characterization of both tasks and users, node activity management and task planning. Moreover, the computing application developed must be easy to manage, so that all operations defining operational and information collection parameters are carried out through an integrated graphic network interface. Lastly, the design should be modular and configurable in such a way that, for each installation, a series of parameters can be defined which enable the operation of the computer application to be configured according to the specific characteristics of each DPC, and where the addition of new modules in order to endow the application with other functions, is a simple process.

3.1. Energy consumption control

Towards the end of the last decade, measures for reducing energy costs in consumer computers began to be analyzed and implemented, above all as regards screens, which represent a significant proportion of total desk top computer energy consumption. An indication of this is the fact that the USA Environmental Protection Agency and the Energy

Department were the impetus behind the EnergyStar label for these units. However, the great leap forward in adopting these energy saving measures came from the generalized mass use of mobile units, above all telephones, PDAs and laptops, which obliged manufacturers to develop low-consumption electrical elements that prolonged battery life [1,4]. Although this development was primarily driven by economics, and did not include an environmental protection perspective, it has nevertheless encouraged the design of processors, memories and other elements with lower consumption rates.

In the wake of these developments, most of the largest computer manufacturers have become involved in Green Computing initiatives, often within the framework of Corporate Social Responsibility strategies, designing more powerful systems with less environmental impact, both in terms of the materials employed and in terms of operational energy requirements. Furthermore, almost all the major companies in the sector (Intel, IBM, Sun, HP, Dell, etc.) have begun to incorporate aspects of rational energy use in their facilities.

Another example of the attention energy saving measures are receiving in the computing sector is the proposal to establish standards which would enable the comparison of systems, not only as regards calculating power but also in terms of the relationship between calculating power and the energy required to achieve it. Thus, for example, SPECpower [8,14] has recently emerged, a performance analysis (benchmark) program which enables the comparison of computer systems as regards both performance and energy consumption.

At present, efforts to reduce the energy consumption of the electronic devices comprising the circuitry are being combined with investment in virtualization strategies which, as mentioned earlier, attempt to reduce the circuitry necessary to optimize use through the implementation of virtual machines. Thus, some work has been done on virtualization systems, such as VMWare, to activate and deactivate virtual machines according to demand.

Clearly, efforts to develop measures to control energy consumption are currently in the initial phase, but they will be of great significance and relevance in the near future. The AECePro project, for example, explores this line of research in depth, and has presented a novel proposal which is perfectly compatible with developments in the energy saving field, both as regards circuitry and software.

3.2. Cluster management tools

Currently, there are a multitude of tools available for the management of computer clusters, such as Rocks, Ganglia, SGE, OpenMosix, (...). The functions which these tools offer concentrate on:

- Help systems for cluster installation, configuration, and maintenance: adding nodes, configuring them, user management, applications control, etc.
- Systems for batch execution of tasks, which permit exclusive use of nodes and balance task load through

a series of functions which normally incorporate a system of queues and a set of task execution planning strategies which can be personalized in each installation.

- System monitoring tools, which provide information regarding the status of each node in the system.

Such tools are developed as middleware which can be installed either as an ad hoc distribution for the operating system (generally Linux), which facilitates installation as it is provided as an integral part of the system, or as software incorporated on top of the operating system. In the later case, the installation and fine-tuning process is much more expensive.

The objective of these tools is two-fold: on the one hand, they attempt to facilitate set-up and maintenance of clusters, and on the other, they help to obtain greater productivity. To date, none of the most widely used distributions for this type of system include modules responsible for minimizing the energy requirements of clusters.

4. Monitoring Project

4.1. Future projects

The Castile and León Supercomputer Center Foundation will be undertaking a project for Integrating Performance Monitoring Systems in the coming months, due to finish at the end of August 2010, comprising the open source application *Nagios* and the HP real time consumption monitor, *Onboard Administrator* (HP is one of the Foundation's many suppliers). The final conclusions of this integration will be implemented through a Plan for Energy Efficiency. This project will entail a significant effort in terms of hours of technical human resources and management, in addition to redeploying the equipment in order to achieve operative integration: system characterization and optimization, and implementation of the plan for energy efficiency.

4.1.1. Integration of Monitoring Systems

Nagios is a widely used *open source* network monitoring system, which monitors the hardware and the software specified, indicating when the behaviour of these is not satisfactory. Among its principle characteristics are monitoring network services and hardware resources (processor loading, disc use, memory, gateway status, etc), independently of the operating system installed, thus facilitating the possibility of remote monitoring via encoded SSL or SSH tunnels, and the possibility of programming specific *plugins* for new systems.

It comprises extremely versatile software which enables practically any system parameter of interest to be examined, and sends a warning to the relevant person when these parameters exceed the margins defined by the administrator.

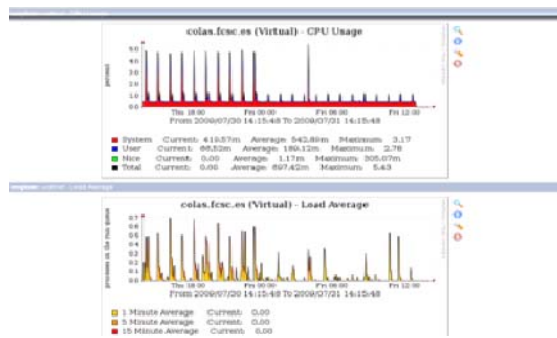


Figure 13. Operational Software (1)

The **Onboard Administrator** for the C7000 chassis is the brain controlling the operation of the HP blade system. This module provides functions such as:

- An Assistant for simple, rapid installation and configuration.
- High availability and security when accessing the blade system infrastructure..
- Security controls for server, network and storage administrators.
- Automation of blade power supply and cooling.
- Temperature and cooling information.

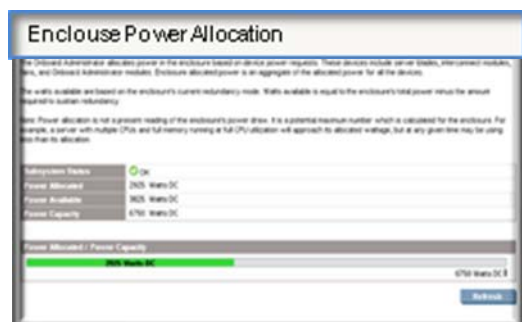


Figure 14. Operational Software (2)

4.1.2. Characterization: An ecosystem called DPC

If we understand a DPC as comprising a series of components which constitute an ecosystem in their own right, it would serve no purpose to have modern, efficient energy transformers and cooling systems if the question of how to reduce the equipment’s energy consumption were not also addressed. It is necessary to understand the behaviour and multiple interactions of a complex system such as a data center, and we will attempt to answer a series of questions:

What can we virtualize? Establish which physical servers have the greatest capacity for specific virtual machines. Monitoring will enable us to select machines which are infrequently used, and which thus become candidates for virtualization. IT resources are rarely shared between different applications, and this can generate fragmented infrastructures and lack of flexibility. We are faced with servers which have a low rate of usage and we do not take advantage of hardware, space or electricity.

A virtualized environment results in the dynamic assignment of resources.

Virtualization also simplifies and speeds up expansion and renovation of installations. The incorporation of resources can be separated from acquisition of hardware.

Another of the benefits of monitoring is that whilst the consolidation of servers based on virtualization can offer many advantages, it can also increase complexity if the environment is not correctly administered. Savings derived from hardware consolidation are lost in increased fixed costs of administration.

Efficient and integrated monitoring addresses questions such as, how can we shut off without losing service quality?, To shut down or disconnect?, and, When does it begin to be economic to shut down the unit (considering the energy peak generated by starting up)?

Efficient and integrated monitoring enables the problem of a reduction in energy efficiency as calculation power increases to be addressed. At a given moment, any additional calculating capacity implies a significant overload in energy terms, causing consumption to shoot up. Characterization of that precise point would enable us to offset consumption with an insignificant loss of calculating power.

HP’s **Onboard Administrator** tool enables a maximum consumption level to be fixed, which cannot be exceeded.

- There are a series of time zones when it would be economic to shut down certain machines. It is of the utmost interest to

know when these are. Some of the more infrequently used servers can be virtualized, and their workload transferred to another, physical, server, making it possible to combine various different applications in one machine. Meanwhile, the original server can be shut down. Later, when activity increases during the day, the servers start up again and the applications are transferred back to them.

- In the calculation nodes, those which are not carrying out, and are unlikely to carry out, tasks, can be shut down. Do not leave them on *stand-by*, which is the strategy currently employed in data centers (AECePro Project).

4.1.3. Plan for Energy Efficiency

The objective of this action plan is to control and reduce energy demand, and to act selectively in terms of energy consumption and supply.

In order to achieve a significant and lasting reduction in energy consumption, it is necessary to develop energy efficient techniques, products and services, and put them into practice.

Once monitoring has been incorporated into the system and the machine has been characterized, a handbook describing the system's energy performance and management protocols needs to be drawn up.

4.2. Viability model and prognosis for sustainability

The actions described here will have continuity over time, and they will become standard operational practice at the center. They will, therefore, have to adapt to future equipment changes and renovation, in addition to new state-of-the-art developments.

The FCSC's mission is to constitute a point of reference in Energy Efficiency over time, and this implies a continuous effort to up-date knowledge in this field, and to monitor all associated new ideas and developments.

The protocols and tools developed will become a reference for other types of centers. For example, in Spain, the criterion of energy efficiency has not previously been considered crucial in the design of supercomputing centers.

We rule out any kind of development which would result in a negative environmental impact, thus ensuring sustainability of the proposed project.

Conclusions

A. General conclusions

A chart of the findings has enabled the identification of the following objectives and proposals:

Economic aspects:

Go from	to
Increase in energy prices	Capacity to fine-tune management of energy costs
Reduction in IT budget	Cost-savings due to efficient energy use
Restrictions to IT growth	Reduction in electricity consumption and cooling (growth inhibitors)

Operational aspects:

Go from	to
High density servers	Increase computation per kilowatt
Soaring electricity and cooling costs	Devote more energy to computation than to physical installations.
Old DPCs	Lengthen the "efficient" life of old DPCs

Environmental aspects:

Go from	to
Social responsibility	Reduction in CO2 emissions
Public image	Improved public image
Professional team morale	Positive impact on professionals from working in an organization which reconciles corporate values with social and personal values.

B. Specific conclusions: specific objectives planned at the Technological Center for Supercomputing in Castile and León.

As regards the *Power Usage Effectiveness –PUE –* commonly found in the 2 industry, that is, the effective combination of industrial technologies and optimization techniques, the present document presents the solution implemented at the Technological Center for Supercomputing in Castile and León, foreseeing a 1,2 result which will situate the center among the first 5 installations world wide due to breaking the 500 Mflops/w barrier.

The energy production mix in Spain generates CO2 emissions of 340 gr/Kwh.

Running at 70% capacity, we predict a consumption of 40.46 Kg of CO2 per hour, 971 Kg of CO2 per day and 354 T of CO2 per year (compared with 562.9 T of CO2 in a 270 Kw installation).

Savings will represent 208.5 t of CO2 per year.

References

1. Francis Zane, Girija Narlikar, Anindya Basu. "CoolCAMs: Power-Efficient TCAMSs for Forwarding Engines". Proc. IEEE INFOCOM 2003, pp. 42–45.
2. Brad Allenby, Darian Unger. "Information Technology Impacts on the U.S. Energy Demand Profile". E-Vision 2000 Conference.
3. Brad Allenby. "Creating Economic, Social and Environmental Value: an Information Infrastructure Perspective". Int. Journal of Environmental Technology and Management. Vol. 7 (5/6) pp. 618–631, 2007.
4. Adrian M. Caulfield, Laura M. Grupp, Steven Swanson. "Gordon: Using Flash Memory to Build Fast, Power-efficient Clusters for Data-intensive Applications". To appear in Proc. 14th Int. Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'09).
5. Neil Rasmussen. "Implementing Energy Efficient Data Center". American Power Conversion White Paper 114, 2006. <http://www.apc.com>.
6. Matthew L. Massie, Brent N. Chun, David E. Culler. "The Ganglia Distributed Monitoring System: Design, Implementation, and Experience". Parallel Computing, vol 30(7), pp. 817–840, 2004.
7. J. A. Gil-Martinez-Abarca, F. Macia-Perez, D. Marcos-Jorquera, V. Gilart-Iglesias. "Wake on LAN over Internet as Web Service". Proc. IEEE Conference on Emerging Technologies and Factory Automation, pp. 1261–1268, 2006.
8. --. Spec Power2008. http://www.spec.org/power_ssj2008
9. --. "Sun Grid Engine (SGE)". <http://docs.sun.com/app/docs/coll/1017.4?l=en&q=Sun+Grid+Engine+6.1>.
10. --. "Intelligent Platform Management Interface Specification v2.0". <http://www.intel.com/design/servers/ipmi/index.htm>.
11. --. "Creating the Green Data Center". <http://www.adc.com>.
12. --. "Creating a Green Data Center to Help Reduce Energy Costs and Gain a Competitive Advantage". IBM Report, 2008. <http://www-935.ibm.com/services/us/cio/outsourcing/gtw03020-usen-01.pdf>.
13. --. "Results of the U.S. Department of Energy and U.S. Environmental Protection Agency's National Data Center". Energy Efficiency Strategy Workshop, 2008. http://www.energystar.gov/ia/partners/prod_development/downloads/energy_eff_data_centers_rec.pdf.
14. --. "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE". The Green Data Center, 2007. <http://www.thegreengrid.org>.
15. --. "Abordar el Reto de la Eficiencia Energética mediante las Tecnologías de la Información y la Comunicación". COM(2008) 241 final. Comunicación de la Comisión al Consejo, al Parlamento Europeo, al Comité Económico y Social Europeo y al Comité de las Regiones.
16. Rajkumar Buyya. High Performance Cluster Computing: Architectures and Systems, Volume I. Prentice-Hall, 1996.