

Energy efficiency in a supercomputing center: a case study

Fernández González, A.¹, Matellán, V.¹, Martínez García, J. M.¹, Lorenzana, J.¹ and López, M.¹

¹ Fundación Centro de Supercomputación de Castilla y León
 Edif. CRAI-TIC, S/N – Campus de Vegazana,
 24071 León (Spain)

Phone/Fax number:+0034 987 293160, e-mail: alvaro.fernandez@scayle.es, vicente.matellan@scayle.es,
jose.martinez@scayle.es, jesus.lorenzana@scayle.es; marivi.lopez@scayle.es

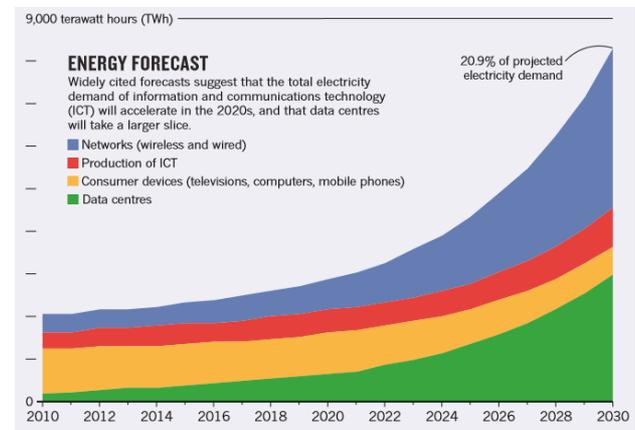
Abstract. The work presents a case study related to the efficient use of energy in the Supercomputing Centre of Castile and Leon (SCAYLE) in the city of León (Spain). In this case study, the location is important to show how weather conditions and an optimal design based in density and free cooling, among others, allows adequate performance of computing infrastructure in terms of energy efficiency, using computing resources in an environment friendly way, maintaining, at the same time, the overall computing performance. In order to analyse the information, the monitoring systems were used to calculate the indicators of performance related to energy consumption. The conclusions of this study shows that, apart from the design of the Data Processing Centre (DPC) an adequate and continuous control of all the components of the Supercomputer are the key for providing the best performance possible in order to obtain very efficient levels of electricity consumption from the computing infrastructure, compared to the total consumption of the Centre.

Key words. Supercomputing, energy efficiency, design, investment, electricity.

1. Introduction

During the last decade, the improvement of energy efficiency has been one of the most important issues in all type of industries. According to a report of National Resources Defense Council (NRDC) [1], U.S. data centers are expected to consume around 140 billion kWh electricity in 2020, with a data center electricity consumption growth expected in 15-20% annually [2], representing almost 1.5% of the world's electricity consumption, that puts ICT's carbon footprint on a par with the aviation industry's emissions from fuel. In graphic 1 an 'expected case' projection is seen and the 'best case' scenario, ICT grows to only 8% of total electricity demand by 2030, rather than to 21%. ICT affects environment in several ways from its production, throughout its use and into its disposal. For this reason, the reduction of electric power consumption and a more eco-friendly way of producing electricity are key to reduce carbon dioxide emissions and their impact on the

environment and global warming [3], while allowing significant cost reductions.



Graphic 1. Projection of electricity consumption. Source: [4]

Supercomputers are great consumers of energy, for maintaining an adequate temperature of the computing infrastructure 24 hours a day, 365 days a year, in fact, in the design of Supercomputers, a key feature is energy efficiency, because the cost of electricity in these infrastructures is very high, representing a growing concern not just for economic reasons but also due to its environmental impact.

For all the reasons mentioned before, power usage levels of top supercomputers represents a burden on system owners and maintainers. In order to analyse this issue, Koomey's law [5] describes an increase of 50% every year in terms of energy efficiency improvement, due to the new generation of hardware platforms, and above all the cooling subsystem, that is currently responsible of 30-60% of final electricity consumption [6]. In next years, exascale level of performance should be based on the infrastructure efficiency related to power limits available in existing computing centers [7].

The interest of the case presented in the study, is based on the description of the industrial infrastructures used for cooling and the reduction of consumption that contributes to an economic improvement in the operation

of a typical Supercomputing Centre, through a reduction of the dependency of the external power supply, reducing emissions that cause the greenhouse effect. Currently, the most important framework for data centers energy efficient is the DCEEF [8] by Green Data Center Alliance (GDCA) that covers five major principles particularly facility design & engineering, information technology, process, governance and finance, providing value for each attribute to represent their energy efficiency ranking in regards to their importance.

The administration of Supercomputers as an infrastructure for research activities, requires a multitude of skills and very focused expertise in High Performance Systems management and an adequate thermal performance [9]. Due to this, a control of all the parameters for a better performance of the technological infrastructure is needed.

In order to analyse the results of the industrial facilities used as auxiliary system to the Supercomputer, it is necessary to make a study of the different alternatives, both before the design of DPC as for the continuous implementation of new technologies to be up-to-date in terms of offering the best resources in every moment.

In the next years, due to issues as growth of the economy, the extended connection of people or the Internet of Things, over 25 billion devices are expected to be connected by 2020 [10].

The present work will focus the study in the energy efficiency in a Supercomputing Centre (SCAYLE), through the description of solutions used currently, operating in the provision of services of calculation and advanced services in the R&D field.

2. The infrastructure of SCAYLE

SCAYLE designed a dense (until 44Kw/rack) and highly innovative DPC committed with the efficiency energy objective, with a minimalist design, adapted to the limited space available (90m²). The DPC is divided into two rooms: the technical room (Fig. 1), that controls the auxiliary systems: electrical panel, fire extinguishing system, systems of uninterruptible power supply, control and distribution of cooling system, etc. and the Supercomputer room, that holds servers (Fig. 2) that contains the computer system based on racks in a cube with closed hot aisle with a exchanged water-air heat between every two racks (Fig. 3), performing water cooling for its greater efficiency and for its good adaptation for a free system indirect cooling. The average PUE of the installation is 1.25.

There is a classification in four categories according to the availability level of the data center, called Tier levels, where I is the lowest with the lowest availability (99.67%) and Tier IV the highest with 99.99% availability. SCAYLE provides a real 24x7x365 infrastructure, making it essential to adapt them, at least, to the level Tier-III of the TIA-942 standard.

Due to the León climatic conditions, the operation of the compressor-free cooling (air-side economizer) is possible most hours of the year, with indirect or direct evaporative cooling providing supplemental cooling during warmer periods.

The features to be considered for the reduction of electrical consumption are the following, according to the experience of the work done in SCAYLE:

1. Hot and cold aisles technique
2. Elimination of “hot points”
3. Rack cooling. Water cooling. “In-row”
4. “Free Cooling”
5. Virtualization
6. Modular design



Fig 1. Technical room: Auxiliary industrial systems



Fig 2. Supercomputer room



Fig 3. Distribution of racks

3. Efficiency energy data

Within a DPC there are two principle sources of energy consumption. One comprises the information systems themselves, including 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, as seen in Figure 4. Clearly, cooling system consumption is closely related to information system consumption. The more servers 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.

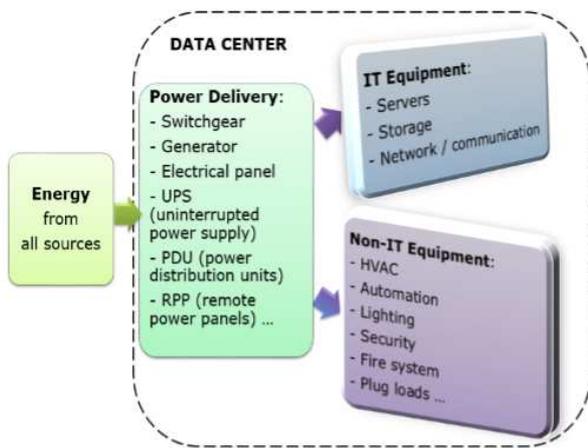


Fig 4. Data center energy flow diagram

Currently, for measuring the efficiency in a DPC, the most used indicator is the PUE (Power Usage Effectiveness), that is a metric developed by the Green Grid [11], widely accepted for measuring energy efficiency, for analysing all the energy consumed by the DPC by that consumed by the servers. PUE formula is:

$$\text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

"Where to measure" will depend on the particular characteristics of each facility. "What" and "how" to measure is the goal of this paper. In its most simplified form, PUE is the ratio between annual energy consumed by the DC and annual energy consumed by the servers. An ideal PUE is 1.0, which means that all energy is consumed by computing equipment and the datacenter infrastructure (such as cooling) takes no energy. As it is almost impossible to reach such ratio, a lot of energy is used by the general equipment of the infrastructure so PUE value is higher than 1.0, being the cooling system of the datacentre one of the most important energy consumer, even in the most optimized cooling systems [12].

In order to improve the management of the PUE ratio, there are recommendations as to take measurements in shorter time periods (monthly, weekly, daily, etc.), whose results may show the evolution of efficiency with time and

give greater understanding of the behavior of the DC to take corrective actions.

In a given time (t), the power used [13] by any system is composed of two parts:

- Pstatic is the fixed part that is due to a gradual loss of energy from charged capacitors;
- Pdynamic is the variable part, resulting from the active usage of components, thus:

$$P(t) = P_{\text{static}} + P_{\text{dynamic}}(t).$$

The energy consumption (E) of a component depends on its power usage (P) over time t. The power is expressed in Watts, and is described as the rate at which the work is performed by the system. Power-related metrics, like max power or average power, usually represent the stress put on components of the studied system. For a given time, interval T1 to T2, with T1 < T2, its energy consumption is given by:

$$E(T_1, T_2) = \int_{T_1}^{T_2} P(t) dt.$$

In this work a measurement method will be treated based on the record of the power demand integrated in 15 min or 1 h periods, due to the ease of using standard measuring equipment to obtain it, the reliability and accuracy of data and usefulness of their results.

The SCAYLE's PUE (Power User Effectiveness) is 1,27 (Total systems power is 140 Kw and the actual electricity consumption is 175 Kwh.), according to the measures made by technicians of SCAYLE following the recommendations found in the literature about where, what and how to measure, and the most adequate time periods for a better interpretation of the results [14].

SCAYLE has installed the transducers necessary for data collection, that provide information of energy consumption, for implementing measures to improve the efficiency of the infrastructure in every moment. Some authors describe the development of models to predict the power profile for the system, that is adapted to the jobs lengths that are running on the supercompyer at the same time [15].

SCAYLE has developed several metrics to control de consumption of energy, in the same line of another entities [16]. The consumption in a period of 12 hours is seen in Fig5. In this graph it can be seen that in most of the time, the instantaneous PUE is 1,2, while in a small percentage of the time the PUE goes up, until peaks of 1, 5. and the consumption in a week in Fig 5.

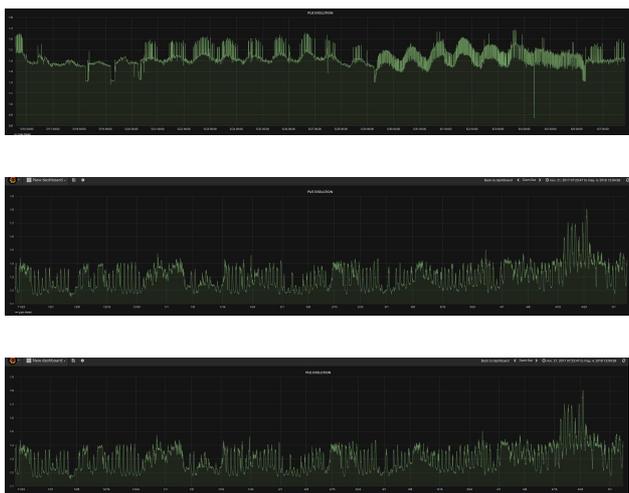


Fig 5. Screens of PUE ratio

As seen in SCAYLE during the analysis of PUE performance, this ratio varies due to some factors:

- System load status: we have seen how are the differences in consumption between servers calculating and without calculating. That is, the equation denominator (ITP) may vary enormously as a function of time.
- Meteorology: in a facility where there is free cooling it is obvious that consumption will vary in function of whether it is being used.
- System configuration parameters: configuration of heat exchangers in the room (in the case of SCAYLE 1 6 APC InRow units RC), water setpoint temperature, temperature of input air setpoint in racks, Vt, etc.

For improving the performance of SCAYLE in terms of the lowest value of PUE possible, all the parameters of the infrastructure are controlled through a panel developed by SCAYLE, in the same line of other projects described in the literature, that shows energy management frameworks offering energy frequency optimisation, energy accounting, and energy/power control [17].

With the information provided SCAYLE efficiency, compared with the Supercomputers with better performance in energy consumption, related to Tflops, is very high as seen on Fig 6, that is the green500 list published every six months. [18]

Green500 List for June 2019

Listed below are the June 2019 The Green500's energy-efficient supercomputers ranked from 1 to 10.

Note: Shaded entries in the table below mean the power data is derived and not measured.

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	469	DGX SaturnV Volta - NVIDIA DGX-1 Volta36, Xeon E5-2698v4 20C 2.2GHz, Infiniband EDR, NVIDIA Tesla V100, Nvidia NVIDIA Corporation United States	22,440	1,070.0	97	15.113
2	1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	10,096	14.719
3	8	AI Bridging Cloud Infrastructure (ABCI) - PRIMERGY CX2570 M4, Xeon Gold 6148 20C 2.4GHz, NVIDIA Tesla V100 SXM2, Infiniband EDR, Fujitsu National Institute of Advanced Industrial Science and Technology (AIST) Japan	391,680	19,880.0	1,649	14.423
4	393	MareNostrum P9 CTE - IBM Power System AC922, IBM POWER9 22C 3.1GHz, Dual-rail Mellanox EDR Infiniband, NVIDIA Tesla V100, IBM Barcelona Supercomputing Center Spain	18,360	1,145.0	81	14.131
5	25	TSUBAME3.0 - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2, HPE	135,828	8,125.0	792	13.704

Fig. 6: Green500 list. June 2019

4. Productivity vs. electric consumption analysis

An evaluation study in relation of productivity based on the number of hours of calculation offered per year for SCAYLE users was made. In case this case, considering that in 2015 and 2018 maintenance works were done, an increase in the electricity consumption is seen on Table 1.

Month	2014	2015	2016	2017	2018
January	43.047	71.136	99.704	123.784	84.403
February	45.350	59.628	73.210	84.696	78.086
March	66.285	75.212	81.058	88.202	82.837
April	63.759	73.569	107.782	99.042	82.405
May	61.501	78.212	93.854	92.824	94.597
June	74.910	74.991	95.426	105.542	98.746
July	71.977	73.921	94.119	104.377	70.417
August	71.823	72.326	101.532	98.675	98.682
September	82.331	82.397	94.914	92.524	98.746
October	76.600	75.896	98.644	73.709	64.051
November	75.390	66.371	82.701	89.805	69.627
December	69.760	70.251	84.563	91.751	71.133
Total	802.733	873.910	1.107.507	1.144.931	993.730

Table 1. Electricity consumption period 2014 – 2018

From the creation of SCAYLE the capacity measured in TFlops was increasing continuously. Currently, The HPC calculation system reaches a power of about 388 TeraFLOPS.

The increase of TFlops described has meant an increase in the calculation hours offered by SCAYLE. In 2015 and 2018 maintenance works were deployed so, the hours of calculation offered in this period were below the capacity of the Centre. In any case, every year, as the renewal of

the infrastructure is done, the hours per Kw consumed is higher, which indicate a better use of energy, as seen on table 2, where is observed that the more increase the hours of calculation, the better is the ratio of hours of calculation vs kw. of energy consumption

Year	Hours of calculation	Ratio h.calc./Kw.
2014	5.833.323	7,27
2015	3.286.532	3,76
2016	9.348.240	8,44
2017	15.183.772	13,26
2018	10.823.845	10,89
Total	44.475.712	9,03

Table 2. Hours of calculation and Ratio hour of calculation vs. Kw

Finally, on table 3 the amount invested in the period studied is seen and also the cost per Kw.

Year	Investment (€)	Ratio invest. (€)/Kw.
2014	294.558	0,37
2015	791.146	0,91
2016	223.119	0,20
2017	96.420	0,08
2018	307.064	0,31
Total	1.712.307	0,34

Table 3. Investment and Ratio investment vs. Kw

5. Conclusions

The first impression of this work is that the phase of design is essential for getting the best results in terms of lower energy consumption. Obviously, the geographical area where the Supercomputing is placed has a great influence, but without an initial study of the best way to reduce the need of energy, no project will success.

The installation of a free-cooling system in Technological Centres with a high demand of energy, has undeniable interest and should be taken in account before investing in a DPC, due to the fast return of these investments that support the decision of purchasing the equipment.

During the last years, many researches were conducted showing cases as cooling systems based on hot liquid, that allows about 10% of possible energy reuse at 50-60°C water temperature that over compensate the performance loss [19]. Most recently, cases as modular design [20], using deep reinforcement learning [21] or the use of microserver high density systems with intelligent power capping mechanisms allows the increasing of energy efficiency of Supercomputing facilities.

Economic reasons will also have great importance for the stimulation of designing equipment's in the way described in this work.

The case presented is running in a Supercomputing Centre with an excellent PUE ratio (1,2) that aims to use best practices for the continuous improvement for the integration of technologies and equipment's to improve energy efficiency, in relation with some studies that shows high power densities and unique cooling system configurations as the key for a better performance of the Supercomputers in terms of reduction of energy consumption [22].

In terms of overall energy use, we can see that centralized data center-based computing in modern facilities is highly efficient.

Through this work we have been able to see that, both in technical terms and on economic terms, it is possible conclude that good practices in the phase of design and in latter phases can be applied to another cases, with major and multiple benefits

References

- [1] P. Delforge, "America's data centers consuming and wasting growing amounts of energy," Nat. Resour. Defense Council Issued Brief, Washington, D C, USA, Tech. Rep., 2014.
- [2] J. Shuja, R. W. Ahmad, A. Gani, A. I. Ahmed, A. Siddiq, K. Nisar and A. Y. Zomaya. Greening emerging IT technologies: Techniques and practices. Journal of Internet Services and Applications, 8(1),9, 2017
- [3] S. Murugesan, "Harnessing green IT: Principles and practices." IT professional10.1 (2008): 24-33
- [4] A. Andrae and T. Edler, On global electricity usage of communication technology: trends to 2030. Challenges, 6(1), 117-157, 2015.
- [5] J. Koomey, S. Berard, M. Sanchez, and H. Wong, "Implications of historical trends in the electrical efficiencyofcomputing,"IEEE Annals of the History of Computing, vol. 33, no.3, pp.46-54, 2011.
- [6] B. Toprak, B. N. Bora and G. N. Güğül, Development of a Simulation Tool to Estimate Electricity Consumption and Determine the Optimum Cooling System for Data Centers. In 2019 IEEE East-West Design & Test Symposium (EWDTS) (pp. 1-6). IEEE, 2019.
- [7] D. Wade, "ASC business plan (NA-ASC-104R-15-Vol.1-Rev.0)," Tech. Rep., 2015.
- [8] Green Data Center Alliance. (2010). Data center energy efficiency framework. Albany, New York: New York State Energy Research & Development Authority.

- [9] T. A. Shedd and R. A. Morell, Cooling 11.6 TFlops (1500 watts) in an office environment. In 2017 33rd Thermal Measurement, Modeling & Management Symposium (SEMI-THERM) (pp. 122-124). IEEE, 2017.
- [10] <http://www.gartner.com/newsroom/id/2636073>
- [11] <https://www.thegreengrid.org/en/resources/library-and-tools/20-PUE%3A-A-Comprehensive-Examination-of-the-Metric>. Retrieved on october 19th 2019.
- [12] M. P. David, M. Iyengar, P. Parida, R. Simons, M. Schultz, M. Gaynes, R. Schmidt, and T. Chainer, "Experimental characterization of an energy efficient chiller-less data center test facility with warm water cooled servers," in 2012 28th Annual IEEE Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM). IEEE, pp. 232–237, 2012.
- [13] A. Oleksiak, T. Ciesielczyk, M. Kierzynka and W. Piatek, Minimising energy costs of data centers using high dense heterogeneous systems and intelligent resource management. In Proceedings of the Ninth International Conference on Future Energy Systems (pp. 499-505). ACM, 2018.
- [14] E. Jaureguialzo, PUE: The Green Grid metric for evaluating the energy efficiency in DC (Data Center). Measurement method using the power demand. In 2011 IEEE 33rd International Telecommunications Energy Conference (INTELEC) (pp. 1-8). IEEE, 2011.
- [15] A. Sirbu and O. Babaoglu, A data-driven approach to modeling power consumption for a hybrid supercomputer. *Concurrency and Computation: Practice and Experience*, 30(9), e4410, 2018.
- [16] F. Peñaherrera and K. Szczepaniak, Development and Application of Metrics for Evaluation of Cumulative Energy Efficiency for IT Devices in Data Centers. In *Cascade Use in Technologies 2018* (pp. 142-153), 2019. Springer Vieweg, Berlin, Heidelberg.
- [17] J. Corbalan and L. Brochard, 2018. EAR: Energy management framework for supercomputers. In *Proc. of IPDPS*.
- [18] <https://www.top500.org/green500/lists/2019/06/>. Retrieved on october 19th 2019.
- [19] S. Zimmermann, I. Meijer, M. K. Tiwari, S. Paredes, B. Michel, and D. Poulidakos, "Aquasar: a hot water cooled data center with direct energy reuse," *Energy*, vol. 43, no. 1, pp. 237–245, 2012.
- [20] R. Roe, Modular design increases energy efficiency: Robert Roe speaks to Bill Thigpen, advanced computing branch chief for the Nasa Advanced Supercomputing Division. *Scientific Computing World*, (164), 14-16, 2019.
- [21] Y. Ran, H. Hu, X. Zhou and Y. Wen, DeepEE: Joint Optimization of Job Scheduling and Cooling Control for Data Center Energy Efficiency Using Deep Reinforcement Learning. In 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS) (pp. 645-655). IEEE, 2019.
- [22] N. Bourassa, W. Johnson, J. Broughton, D. M. Carter, S. Joy, R. Vitti and P. Seto, Operational data analytics: Optimizing the national energy research scientific computing center cooling systems. In Proceedings of the 48th International Conference on Parallel Processing: Workshops (p. 5). ACM, 2019.