

System software for energy management in HPC



The main goal of this group is to provide system software for energy management in HPC data centers. The main project of our team is EAR, an energy management framework developed in the BSC-Lenovo project.

EAR is an energy management framework optimizing the energy and efficiency of a cluster of interconnected nodes. To improve the energy of the cluster, EAR provides energy control, accounting, monitoring and optimization of both the applications running on the cluster and of the overall global cluster including global energy capping which will ensure that the maximum energy values defined for the cluster are not reached.

At EAR's core is a monitoring tool which gathers data on the nodes and on the applications running on the cluster. Therefore, on top of optimizing the energy consumed by the applications running on the cluster and the overall global cluster, EAR reports system and application information.

The system information collected by EAR (called the system signature) reports the performance of the major components of each node (cpu, memory). It is used to optimize the cluster energy but also to report which components are not working to the level expected like, for example, the memory DIMMs in a node are not providing the bandwidth they should, or a cpu in a node is not running at the expected frequency. Therefore, EAR makes sure the performance efficiency of the cluster is kept to its maximum.

The application information collected by EAR (called the application signature) reports basic performance metrics of the application. It is used to optimize the application energy but it can also be used to determine that one application could perform better due to, for example, a high memory bandwidth or a low percentage of AVX512 instruction. Therefore, EAR helps to a better utilization of the system.

EAR is also robust and reliable as it is operational on SuperMUC-NG at Leibniz Supercomputing Centre (LRZ) in Garching near Munich since August 8 2019 (<https://doku.lrz.de/display/PUBLIC/SuperMUC-NG>) running on a cluster of 6480 nodes helping LRZ to meet its energy goals. At LRZ, EAR is transparently used through a SLURM plugin (<https://doku.lrz.de/display/PUBLIC/Energy+Aware+Runtime>).

Objectives

- System software for energy management including

- Simple and accurate job energy accounting
- Automatic energy policies
- Node power and performance monitoring
- Transparent and lightweight runtimes (EARL) for energy efficient applications
- Application and system analysis of energy and performance metrics

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 1 jun 2025 - 16:57): <https://www.bsc.es/ca/discover-bsc/organisation/research-departments/system-software-energy-management-hpc-0>