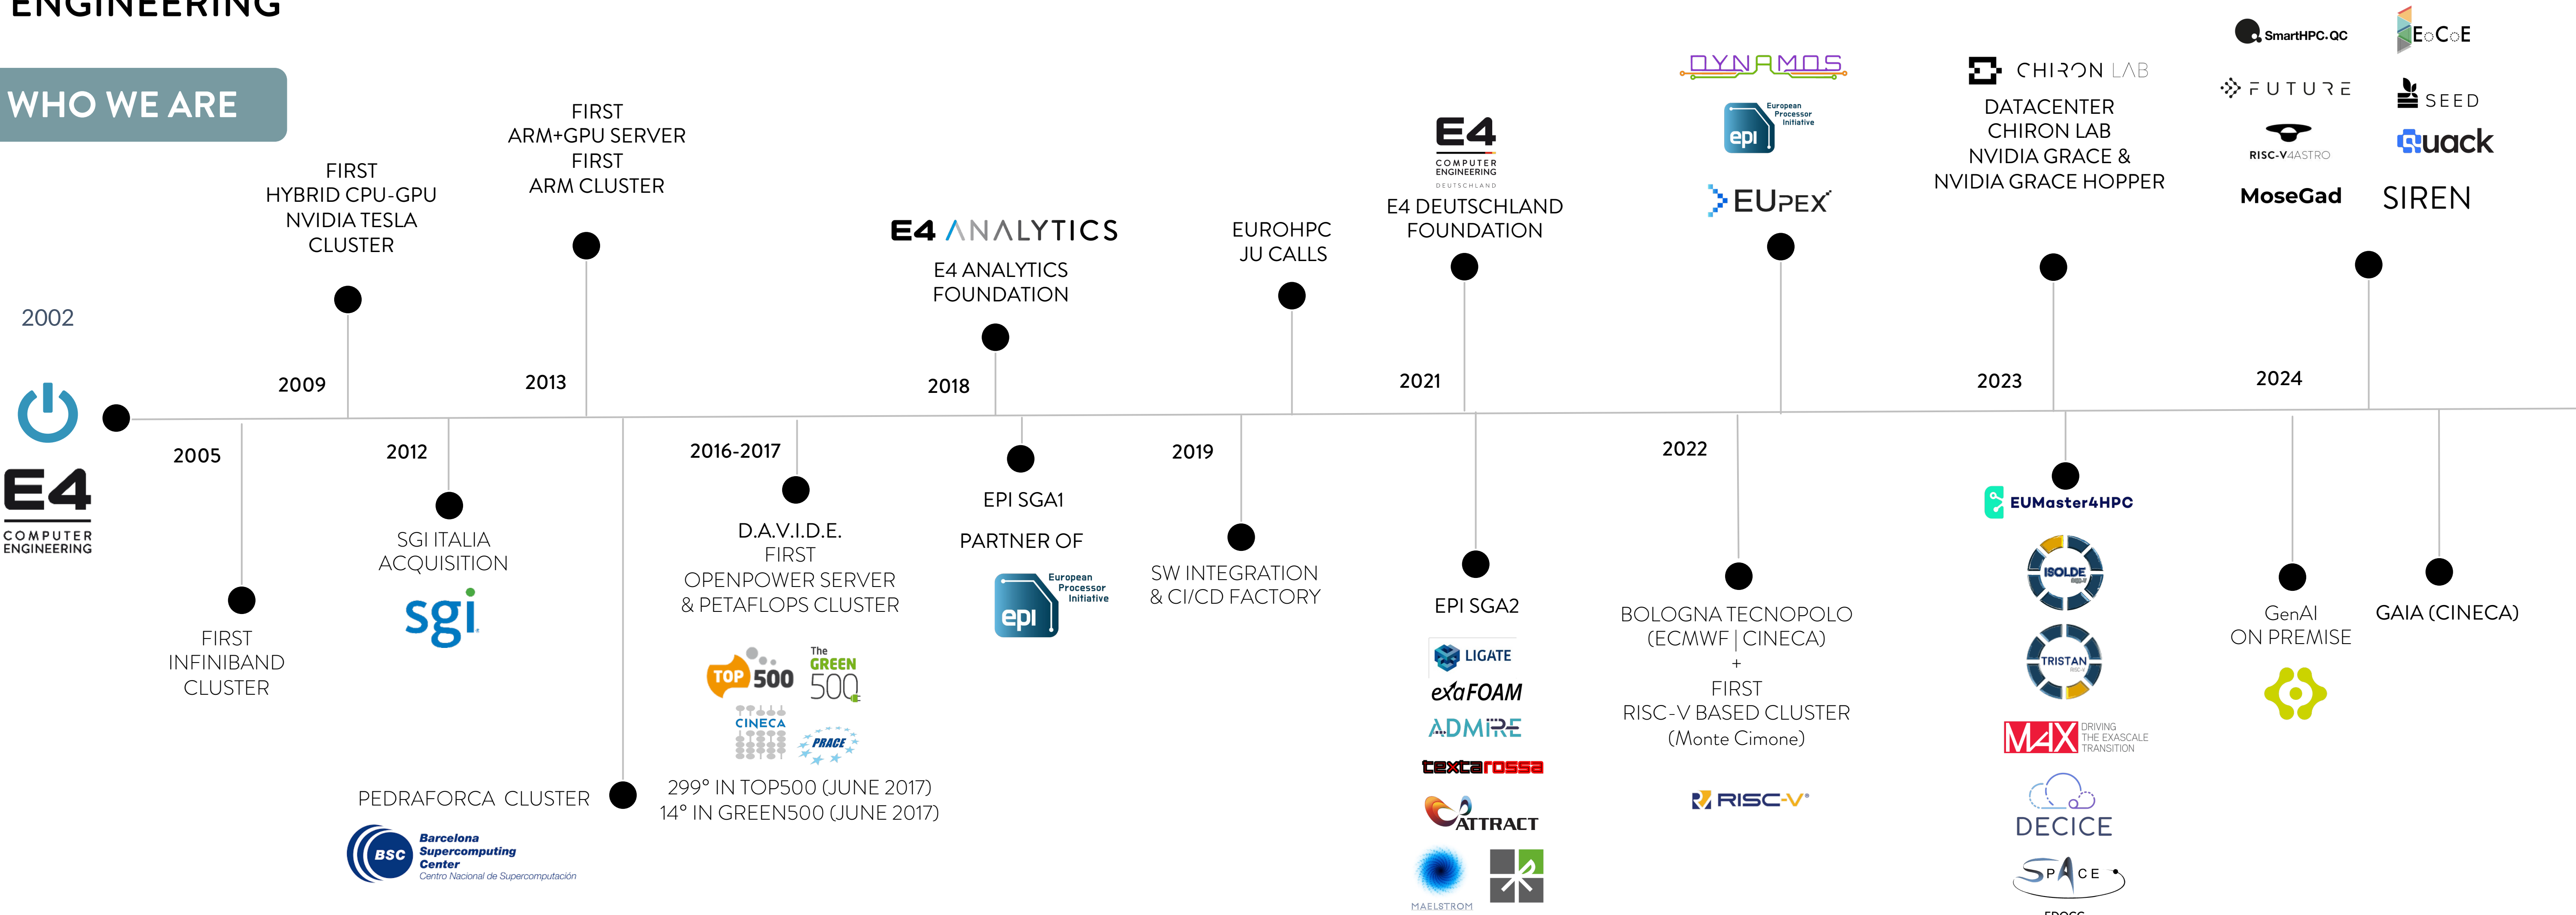



WHO WE ARE



E4 is one of the main HPC service integrator in Europe; we work with all the main hardware vendors for all the components category and our job is codesign with our customers and deploy the infrastructures that best serve their needs. Most of the clusters are used in hard sciences or engineering simulations, in academia or enterprises from pharma to aero space.

E4 IN NUMBERS

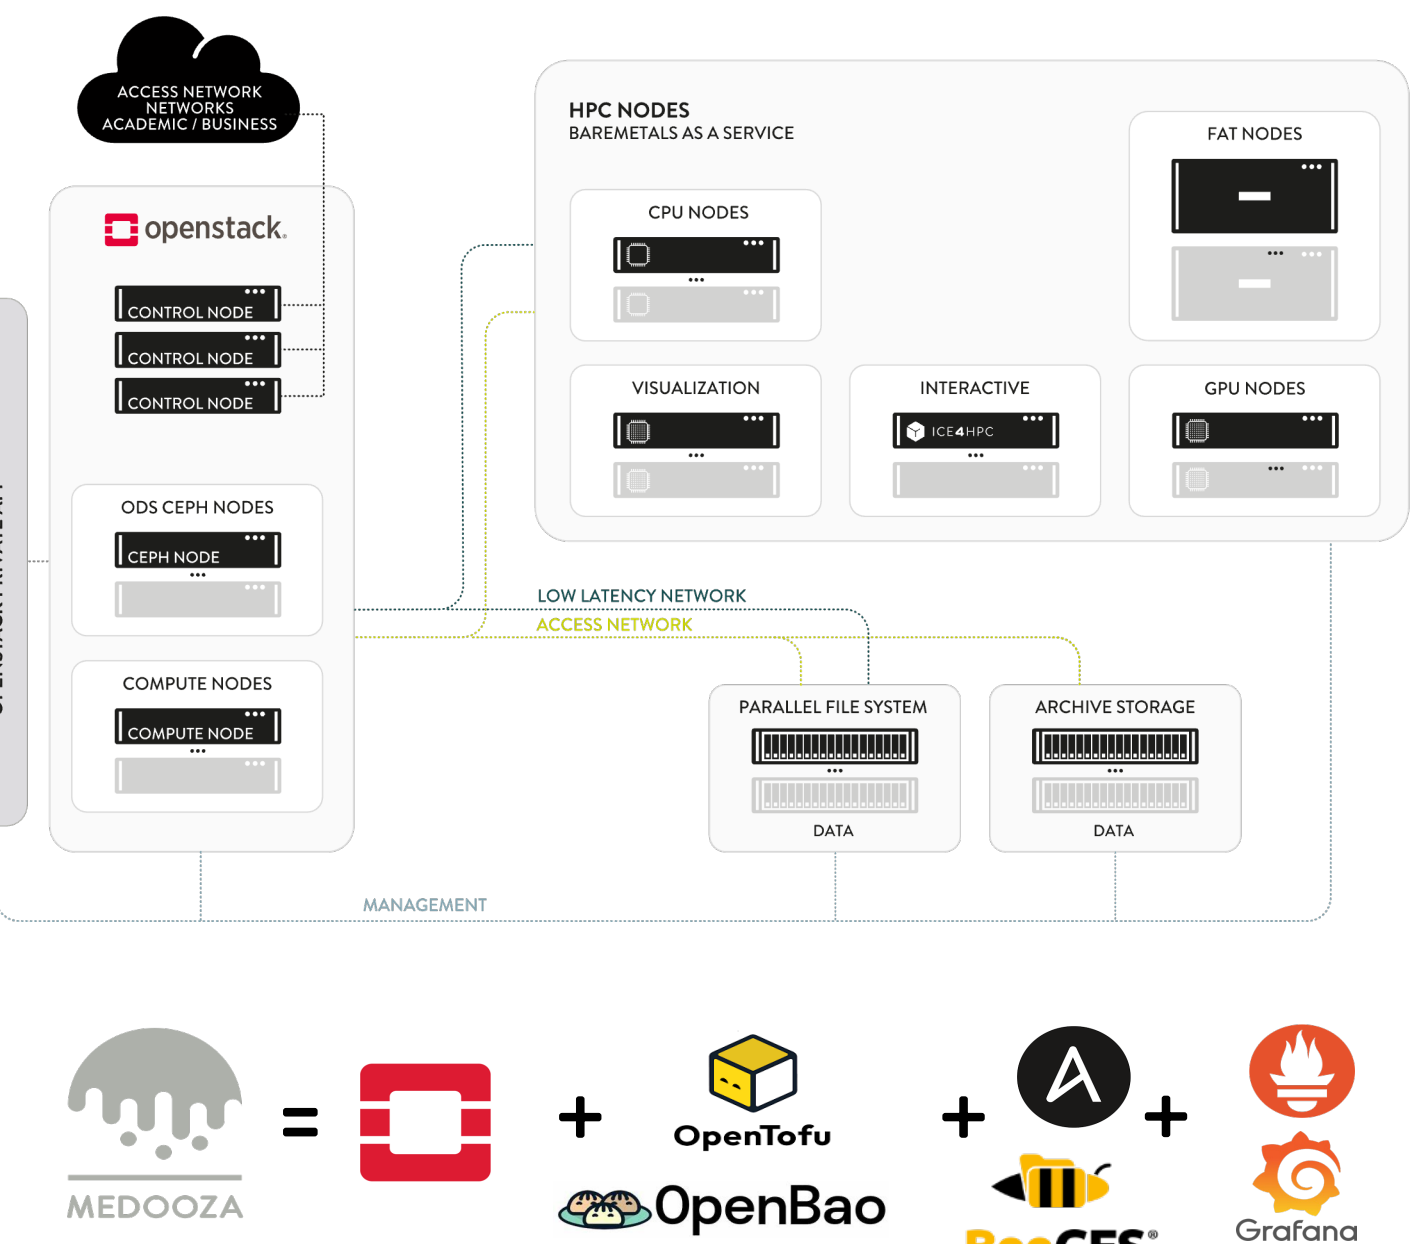


TECH FACTORY


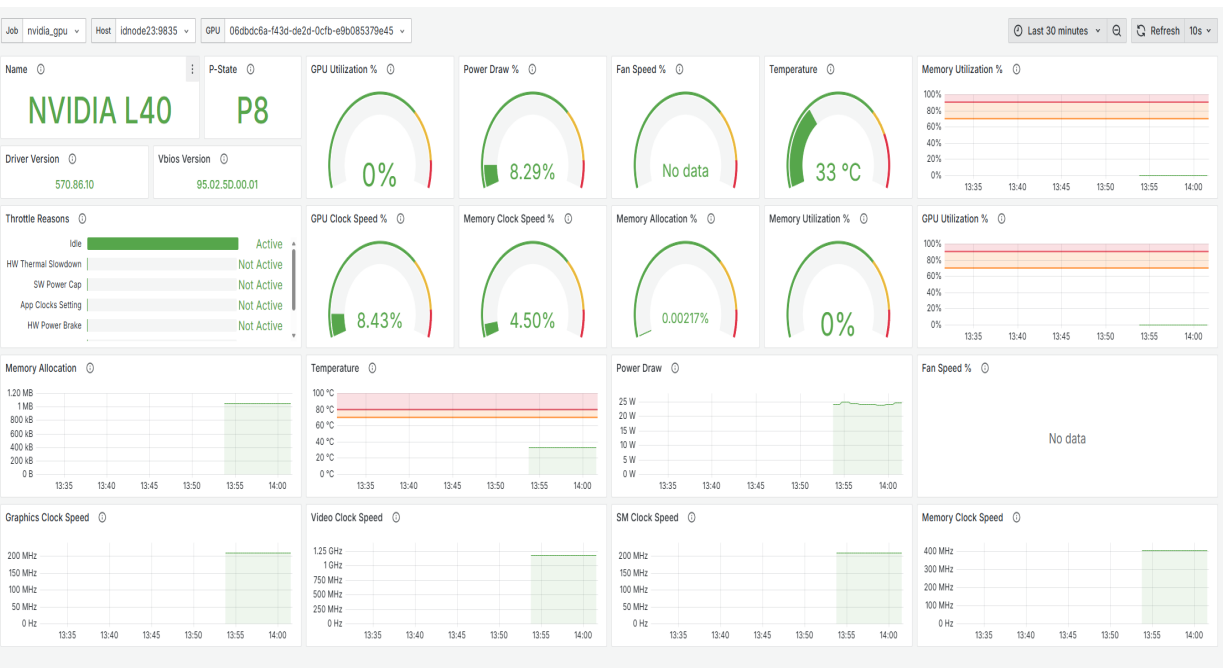


Integration Facility where our technicians build servers or storage systems.
Burn In Room to improve E4 systems reliability with at least 72 hours of test that involves all components. **R&D Lab** with 6 standard racks with heterogeneous systems, 100kW, remote access available on demand to perform benchmarking, co-design, prototyping.

CLUSTER MANAGEMENT



Medooza: General architecture and set of tools to manage and deploy HPC and private cloud. Use Iac tools to create resources and manage secrets, etc.



FLAGSHIP PROJECTS



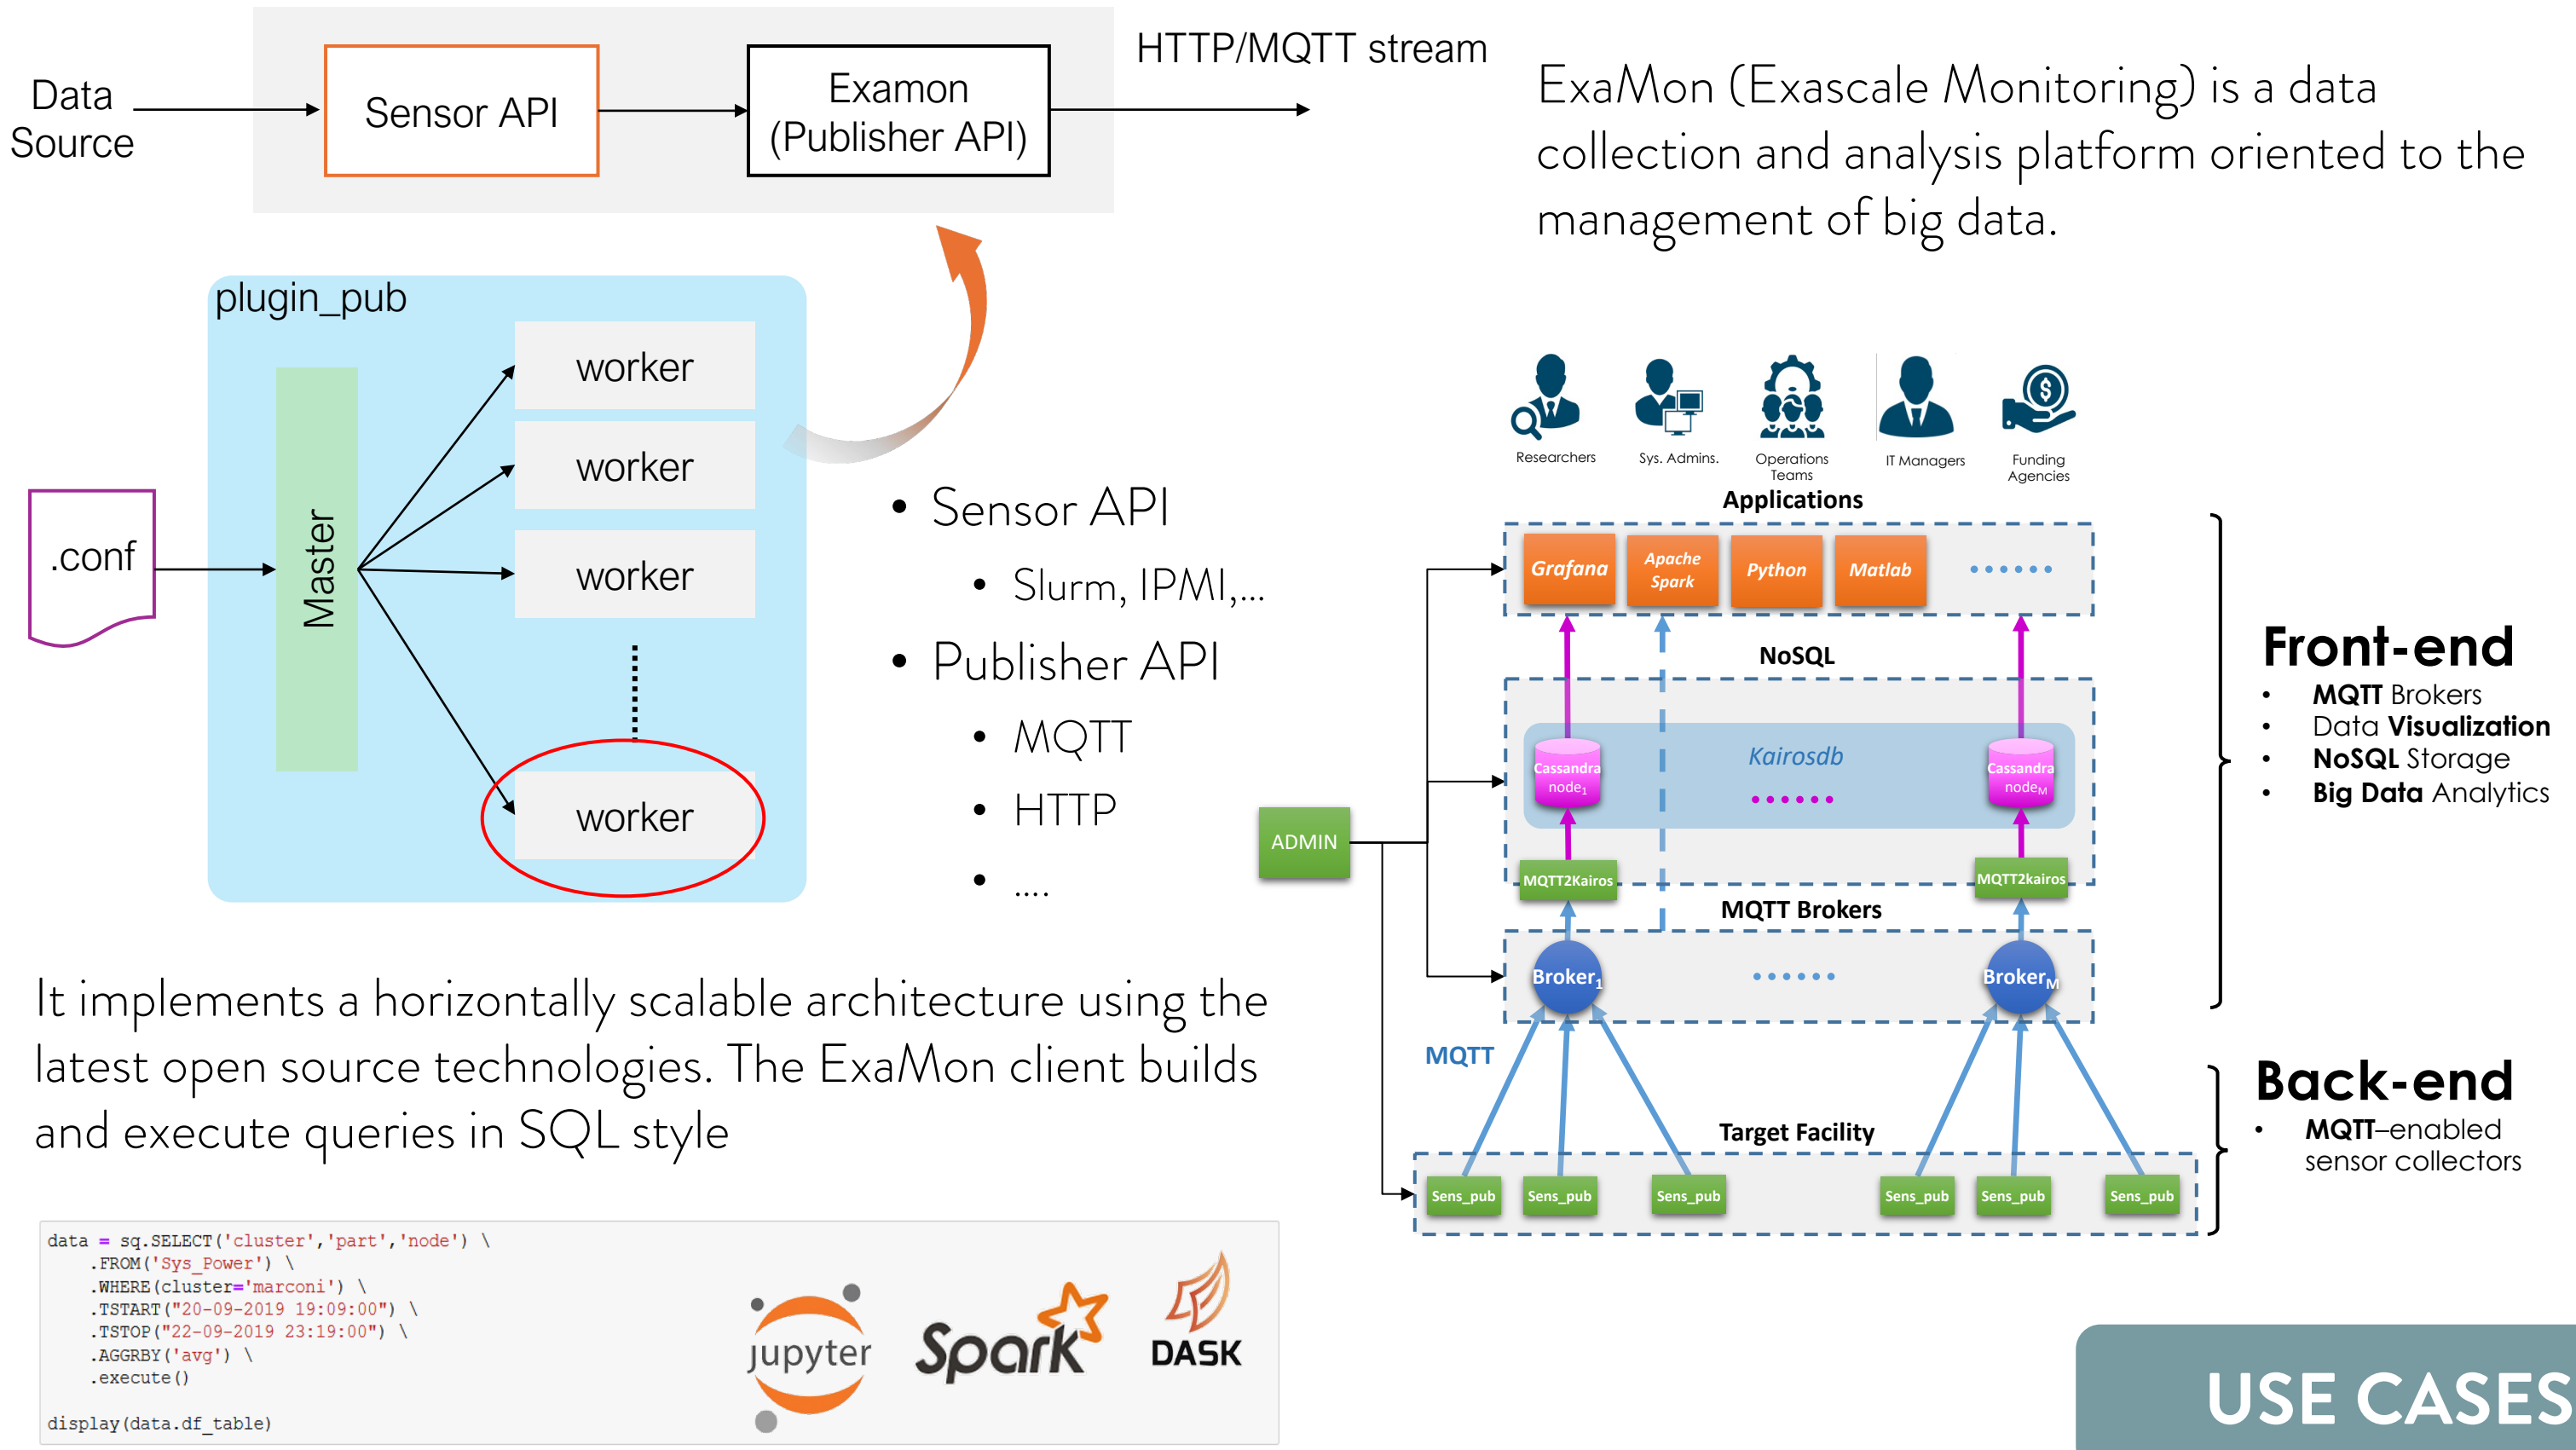
ASTRI@Mount Teide, Tenerife: ICT infrastructure for the pilot project of the Cherenkov telescope array.

Franklin@IIT, Raise@IIT, Genoa: HPC and Private cloud clusters based on Medooza with complete SW stack.

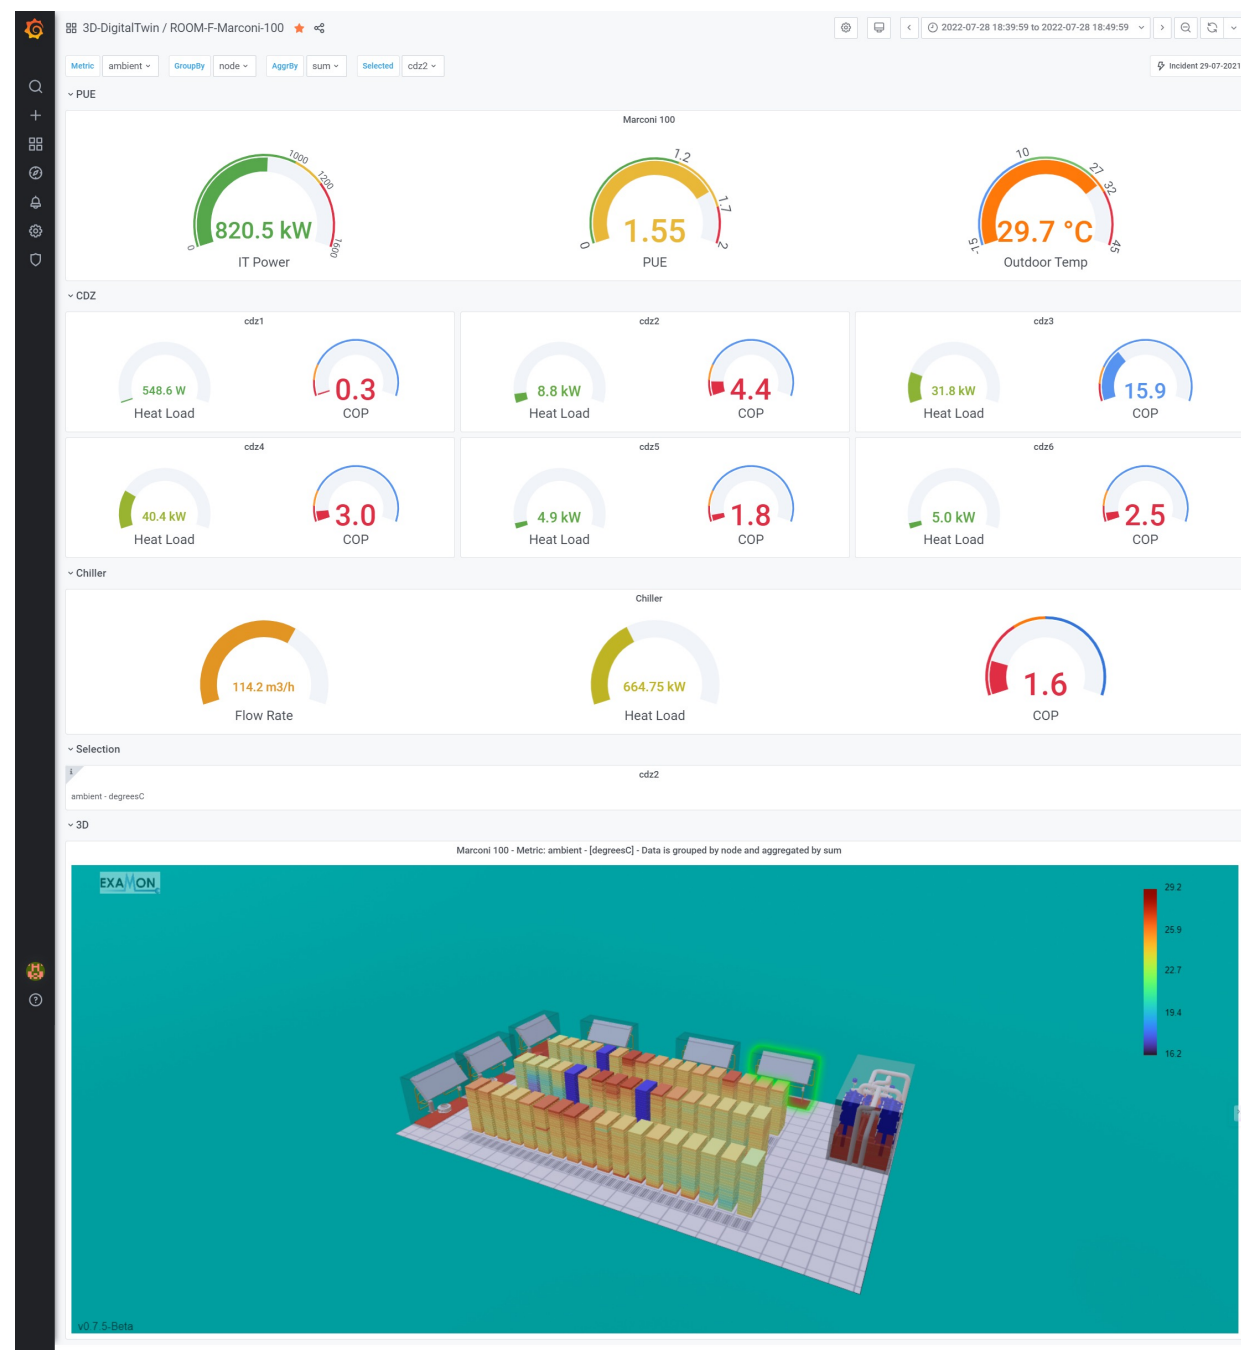


Gaia@CINECA, Bologna: Expansion of Galileo100 called Gaia: cloud for public research in CINECA.

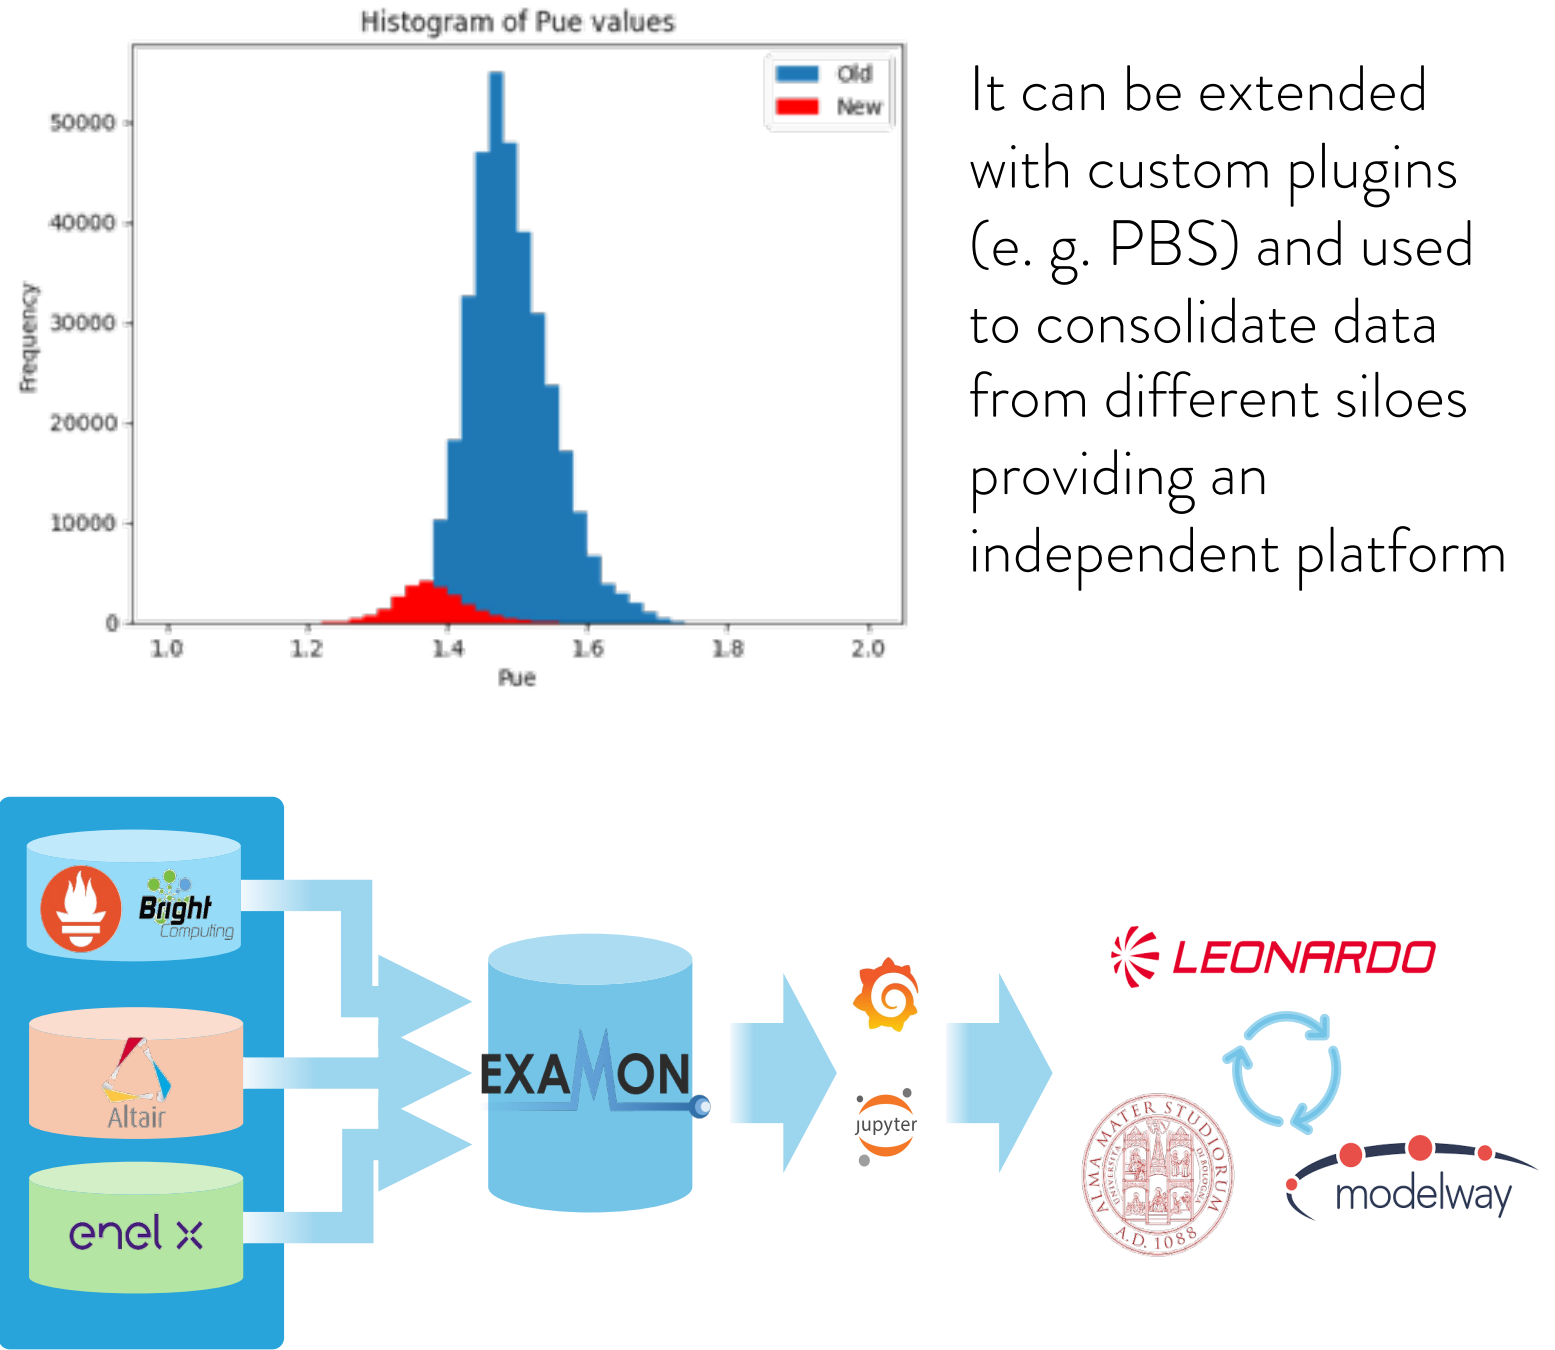
EXAMON PROJECT




Cluster digital twin: the efficiency curves obtained from historical data can be analyzed to define the optimal operating point of the devices as a function of load, temperature, and humidity.



Thanks to the immediate feedback provided by the dashboards, the operators were able to set the individual set points of the devices optimally.
PUE reduction of 8%



SUMMARY AND PERSPECTIVES



ExaMon can be used to collect data per job (type) and per user and the data can be modeled with AI methods to provide predictive energy usage (scheduler hooks and policies can be optimized) and predictive anomaly detection (preventing costly restarts and allowing less checkpointing e.g. for training jobs).

| Job Energy By User | | |
|--------------------|--------------|-----------------------|
| user_name | data_quality | total_energy_Wh (sum) |
| eboella | 100 | 25.10 kWh |
| ggiaquinto | 100 | 5.52 kWh |
| faleotti | 100 | 3.95 kWh |
| odebei | 100 | 1.63 kWh |
| dfuzzati | 100 | 702.58 Wh |

Normal LLM models cannot be used for advanced AD but alternatives exists, see Molan et al. (<https://doi.org/10.1016/j.future.2024.06.032>). This would allow to create operational chatbots for HPC operators see J. A. Khan, M. Molan, e A. Bartolini, «EXASAGE: The first Data Center Operational Data Co-pilot»