# SC17: EE HPC WG Workshop

# Cineca Site Update

Mirko Cestari
Carlo Cavazzoni

| Logical Name | D.A.V.I.D.E. (August 2017) |
|---|---|
| Model | E4 Cluster<br>Open Rack |
| Architecture | OpenPower NVIDIA NVLink |
| Processor | OpenPower 8<br>NVIDIA Tesla P100 SXM2 |
| # of core | - |
| # of node | 45 x (2 Power8 + 4 Tesla P100) |
| # of rack | - |
| RAM per node | - |
| Interconnection | Mellanox EDR |
| Operating System | GNU/Linux |
| Total Power | - |
| Peak Performance | ~ 1 Pflops |

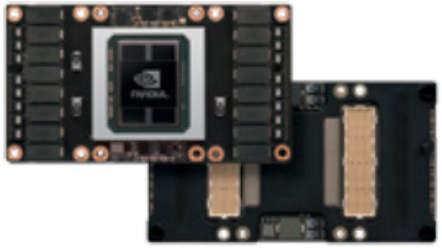

## SPECIFICATIONS

| | |
|---|---|
| GPU Architecture | **NVIDIA Pascal** |
| NVIDIA CUDA® Cores | **3584** |
| Double-Precision Performance | **5.3 TeraFLOPS** |
| Single-Precision Performance | **10.6 TeraFLOPS** |
| Half-Precision Performance | **21.2 TeraFLOPS** |
| GPU Memory | **16 GB CoWoS HBM2** |
| Memory Bandwidth | **732 GB/s** |
| Interconnect | **NVIDIA NVLink** |
| Max Power Consumption | **300 W** |
| ECC | **Native support with no capacity or performance overhead** |
| Thermal Solution | **Passive** |
| Form Factor | **SXM2** |
| Compute APIs | **NVIDIA CUDA, DirectCompute, OpenCL™, OpenACC** |

TeraFLOPS measurements with NVIDIA GPU Boost™ technology

# DAVIDE dedicated SoC: Energy and power gateway (EG)

## Computing node

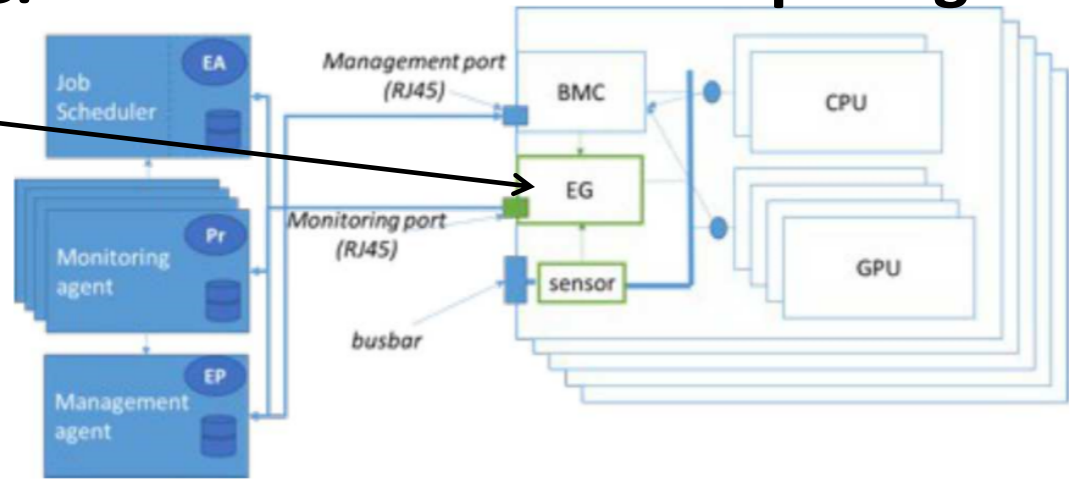


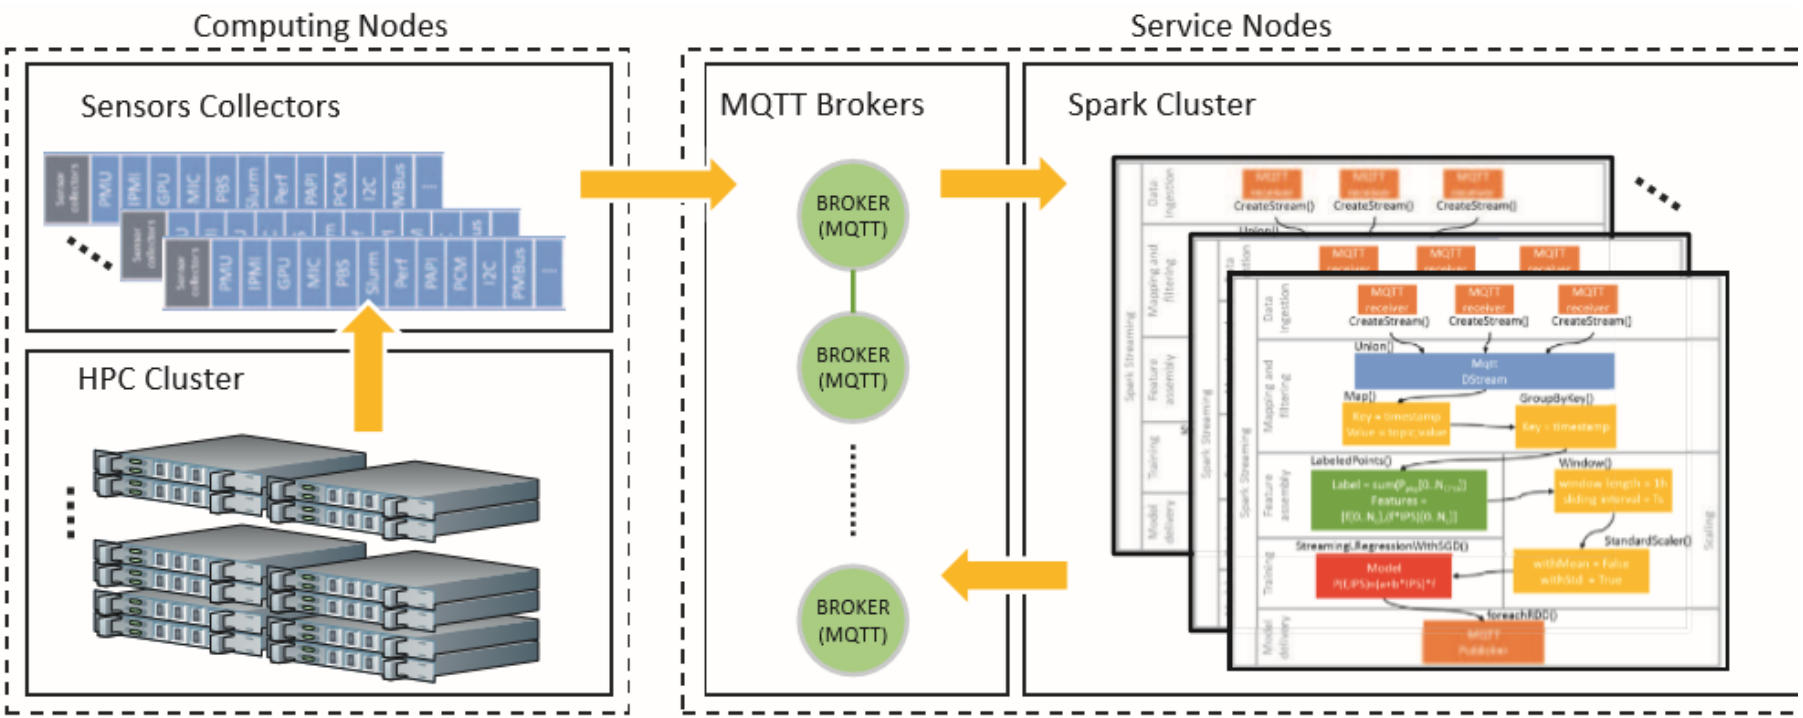

# Data Room: Thermal model

- model each macro component (chillers, system,…) as a sink or source of heating
- Goal: coordinated cooling strategy, including job dispatching
- Needs to be validated
  - Room sensors (map)

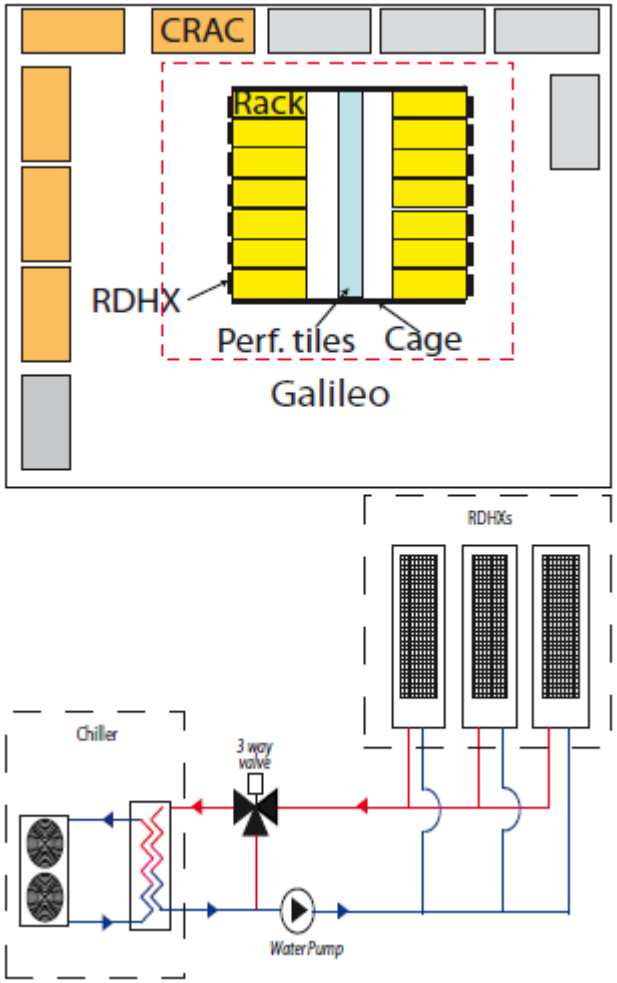


A. Bartolini group,
University of Bologna

# Temperature Data Acquisition

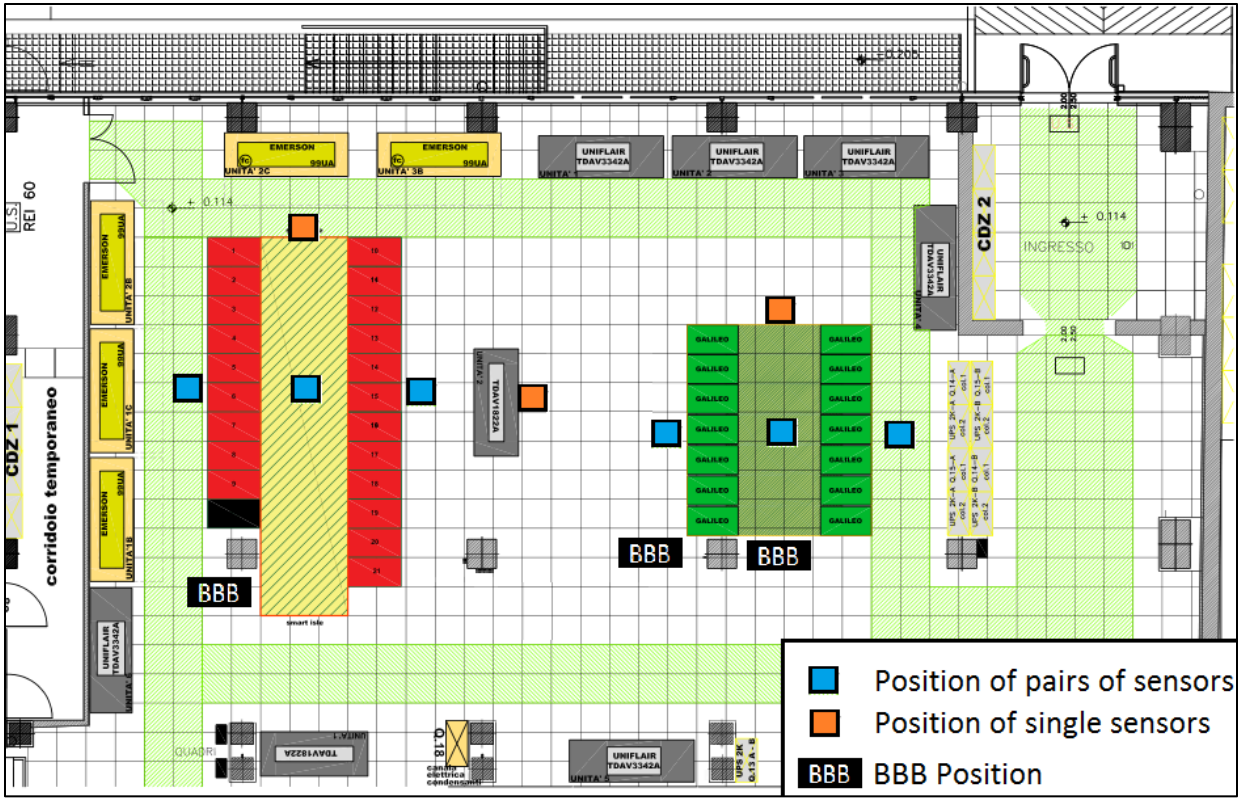


sensors

+

collectors

MQTT

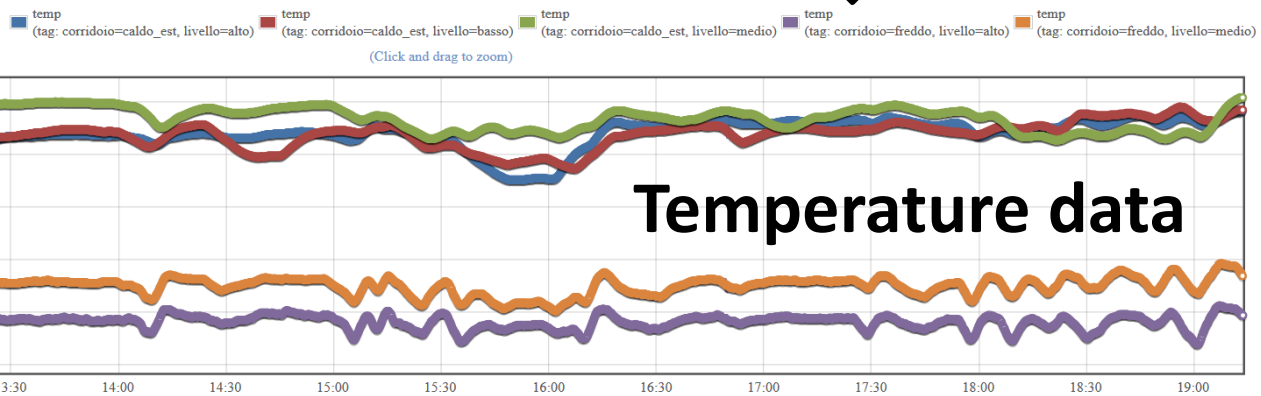
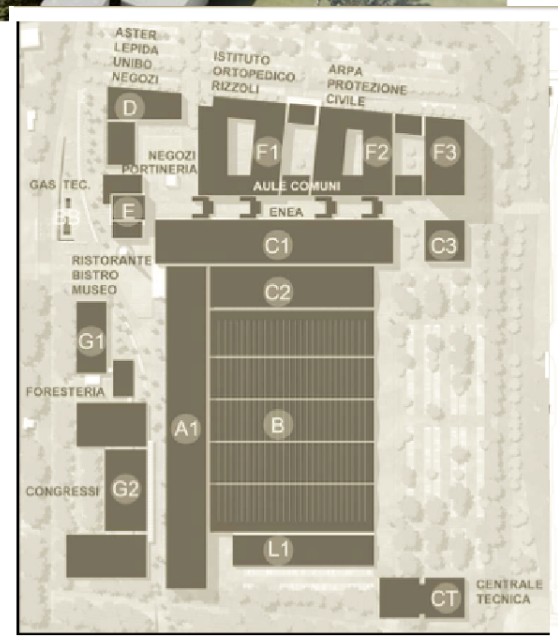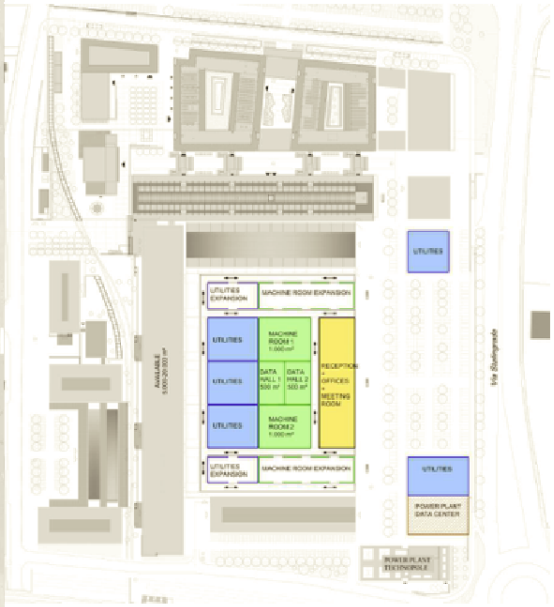
# ECMWF location