



Contribution ID: 61

Type: Oral

## The new GPU-based HPC cluster at ReCaS-Bari

*Tuesday 25 October 2022 15:10 (20 minutes)*

The ReCaS-Bari datacenter enriches its service portfolio providing a new HPC/GPU cluster for Bari University and INFN users. This new service is the best solution for complex applications requiring a massively parallel processing architecture. The cluster is equipped with cutting edge Nvidia GPUs, like V100 and A100, suitable for those applications able to use all the available parallel hardware. Artificial intelligence, complex model simulation (weather and earthquake forecasts, molecular dynamics and galaxy formation) and all high precision floating-point based applications are possible candidates to be executed on the new service. The cluster is composed of 10 machines with a total computing resource equals to 1755 cores, 13.7 TB RAM, 55 TB local disk and 38 high performance GPUs (18 Nvidia A100 and 20 Nvidia V100). Each node can access the ReCaS-Bari distributed storage based on GPFS equals to 8.3 PB. Applications are executed only within Docker containers, conferring to the HPC/GPU cluster features like easy application configuration and execution, reliability, flexibility and security. Currently, users are able to choose among different ready-to-use services like remote IDEs (Jupyter Notebook and RStudio), by which execute GPU based applications, or a job orchestration to whom submit complex workflow represented as DAG (Directed Acyclic Graphs). The user service portfolio is in evolution. If the provided user services do not cover the user needs, user-defined Docker containers can be executed on the Cluster. Long running services and job submission are managed with Marathon and Chronos respectively, two frameworks running along with Apache Mesos. These three tools add high availability, fault tolerant and security additional to the native capacity to manage all compute resources and user requests. The implemented technological solution allows users to continue to access their own data both from HTC cluster (based on HTCondor) and from HPC/GPU Cluster, based on Mesos.

The first phase, where local beta-testers used the cluster, concluded successfully. The service is now ready to join the national INFN-Cloud federation. Leveraging the INDIGO PaaS orchestrator, provides multiple ready-to-used frameworks and services (ML\_INFNO, Apache Spark, JupyterLab, ...), a stable and secure authentication layer, a simple web dashboard that can be used to deploy services on top of and an heterogeneous set of resources. The evolution of the service, where a performance evaluation of Kubernetes as replacement of Apache Mesos, is in the pipeline.

In this contribution will be presented and discussed resources and technological solutions related to the HPC/GPU Cluster in the ReCaS-Bari data center and the most important applications running on the cluster.

### Significance

### References

### Experiment context, if any

**Primary author:** VINO, Gioacchino (INFN Bari (IT))

**Co-authors:** ITALIANO, Alessandro (INFN - National Institute for Nuclear Physics); ELIA, Domenico (INFN Bari); DONVITO, Giacinto (Universita e INFN, Bari (IT)); ANTONACCI, Marica (INFN)

**Presenter:** VINO, Gioacchino (INFN Bari (IT))

**Session Classification:** Track 1: Computing Technology for Physics Research

**Track Classification:** Track 1: Computing Technology for Physics Research