

ReCaS GPU Cluster

Hardware Facility:

- Nodes: 10
- GPUs: 38 (V100 and A100 Nvidia GPU)
- Cores: 1755
- RAM: 13.7 TB
- Local Storage: 55 TB (SSD/HDD)
- Parallel File System: ReCaS storage based on IBM GPFS (8.4PB)
- Bandwidth between nodes: 10 Gbps



ReCaS GPU Cluster: Portfolio

- **Ready-to-use services:**

- Interactive remote GPU-based IDE services:

- Jupyter

- “web service for interactive computing across all programming languages”


- Rstudio

- “An integrated development environment for R”

- Job Scheduler:

- Support to GPU-based workflows represented as Directed Acyclic Graphs (DAG)




- **User-defined services**



ReCaS GPU Cluster: Portfolio

Jupyter Notebook remote IDE

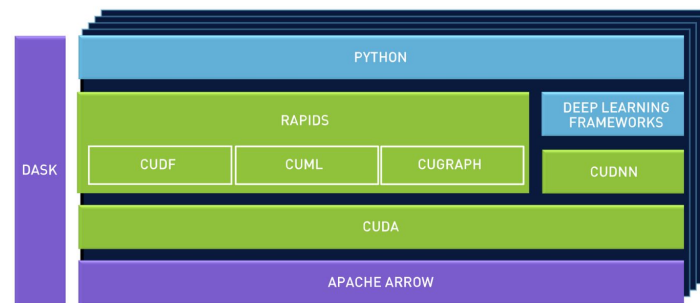
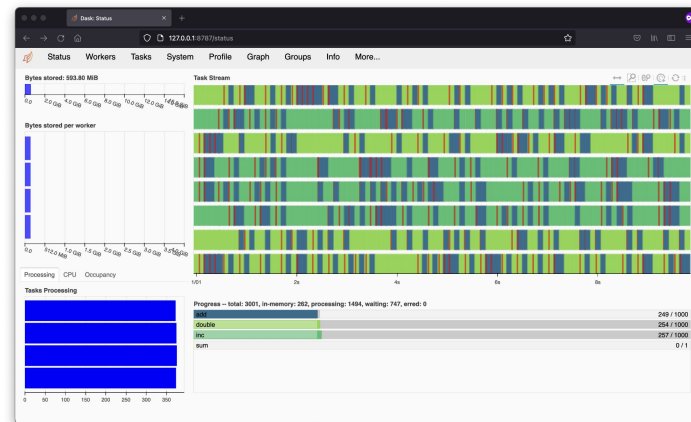
- After authentication, users have access to their home directory in the ReCaS distributed storage (GPFS)
- Users can immediately create a new Python3 script
- The Jupyter IDE (Integrated Development Environment) will be available and users can already write code and execute it
- Python modules can be installed directly within the code



ReCaS GPU Cluster: Portfolio

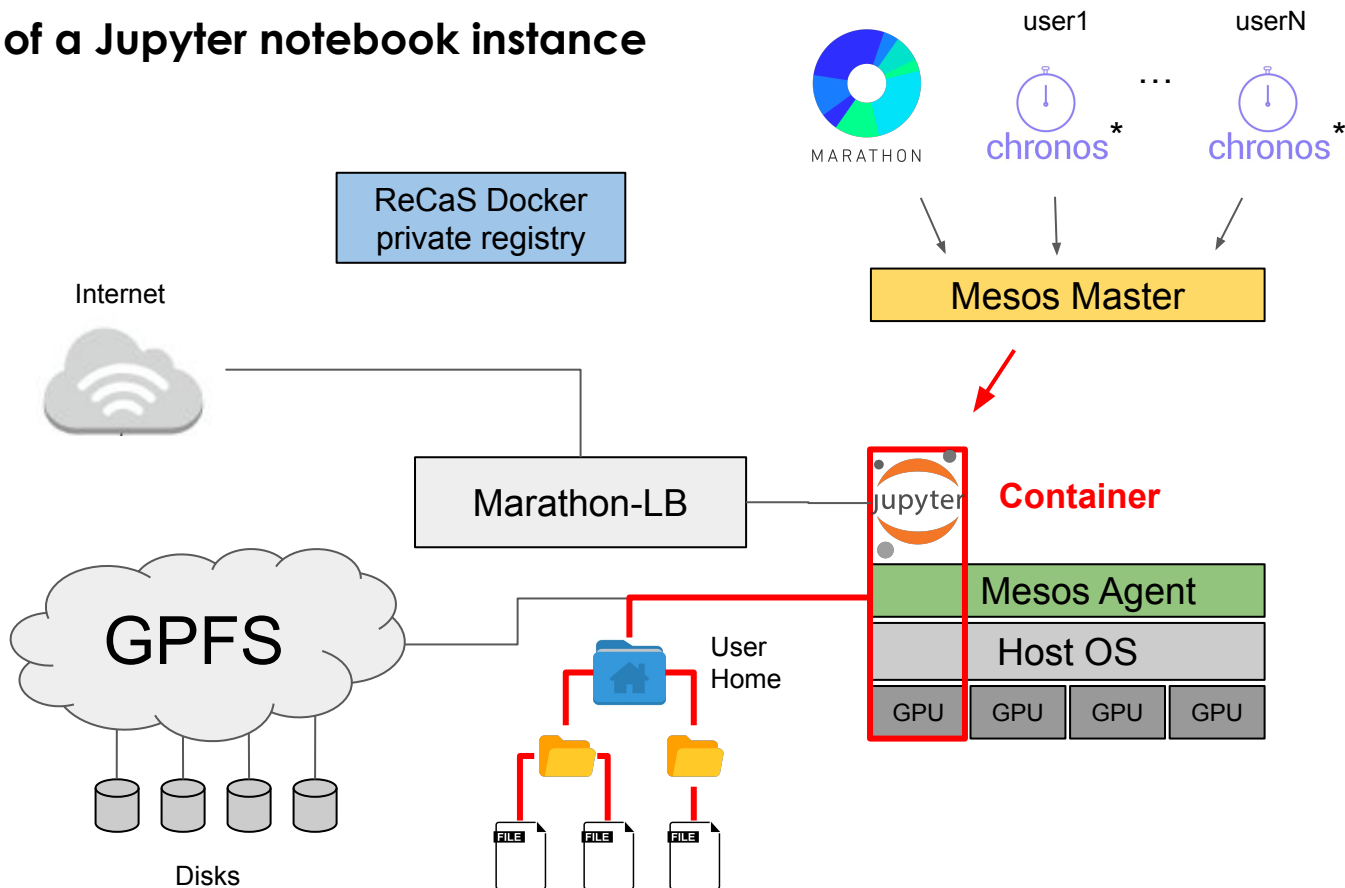
Dask on GPU with Jupyter Lab

- Speed up the single core pandas/numpy algorithms
- Parallel programming library for Python
- Dask copies the pandas and numpy API
- Scales data libraries like numpy, pandas, scikit-learn,..
- GPU optimized pandas DataFrame and numpy Array
Cuda libraries to speed-up the processing using the
GPU Nvidia



How ReCaS GPU Cluster works

Deployment of a Jupyter notebook instance



ReCaS GPU Cluster: Users

Users

- Bioinformatics
 - CNR/IGB
 - Department of Veterinary Medicine
 - Molecular Biology at University of Bari
- AI/ML algorithms
- National and European projects
- Mathematics