Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

UC2.2.2 - Quasi interactive analysis of big data with high throughput

Tommaso Diotalevi (UniBO), Francesco G. Gravili (UniSalento)

WP2 Meeting - 22 July 2025

Milestone 10 - Final Report

- The final version of the intermediate report for MS10 (March 2025) is available in the workflow document.
  - There were no intermediate KPIs associated to this milestone, since all the remaining objectives have been set for the end of project (MS10 final report).

- The final report of MS10 is expected in mid-September 2025, with the finalisation of the KPIs.

- MS11 timeframe will be dedicated to scale testing the WP5 platform and eventual fallback solutions.

Conference opportunities for our flagship

- Contribution given at:
  - ISGC 2025:
    - "*Quasi interactive analysis of High Energy Physics big data with high throughput*" (oral contribution) by *T.Diotalevi* and *F.G.Gravili*
    - *"Distributing the Simulated Annealing workload for Quantum Unfolding in HEP"* (oral contribution) by *S.Gasperini* (Unibo), *T.Diotalevi* and *M.Lorusso* (Unibo/CNAF)

News from WP5

- Total of 70 worker nodes @INFN - NAPLES, shared with other projects, each with:
  - 502GB RAM;
  - 96 core in total, competing with other users;

- Access "grid-like" to the cluster (HTCondor), with CVMFS available (cern.ch & infn.it);

- Not all the cluster is currently active, due to cooling issues hopefully fixed before summer break. Nodes will progressly be enabled, also based on the usage;

For the first phase of scale tests, a subset of such resources have been allocated "privately" for our flagship:

- 5 full worker nodes, extendable in case of need.

News from WP5

Images from T.Tedeschi's presentation

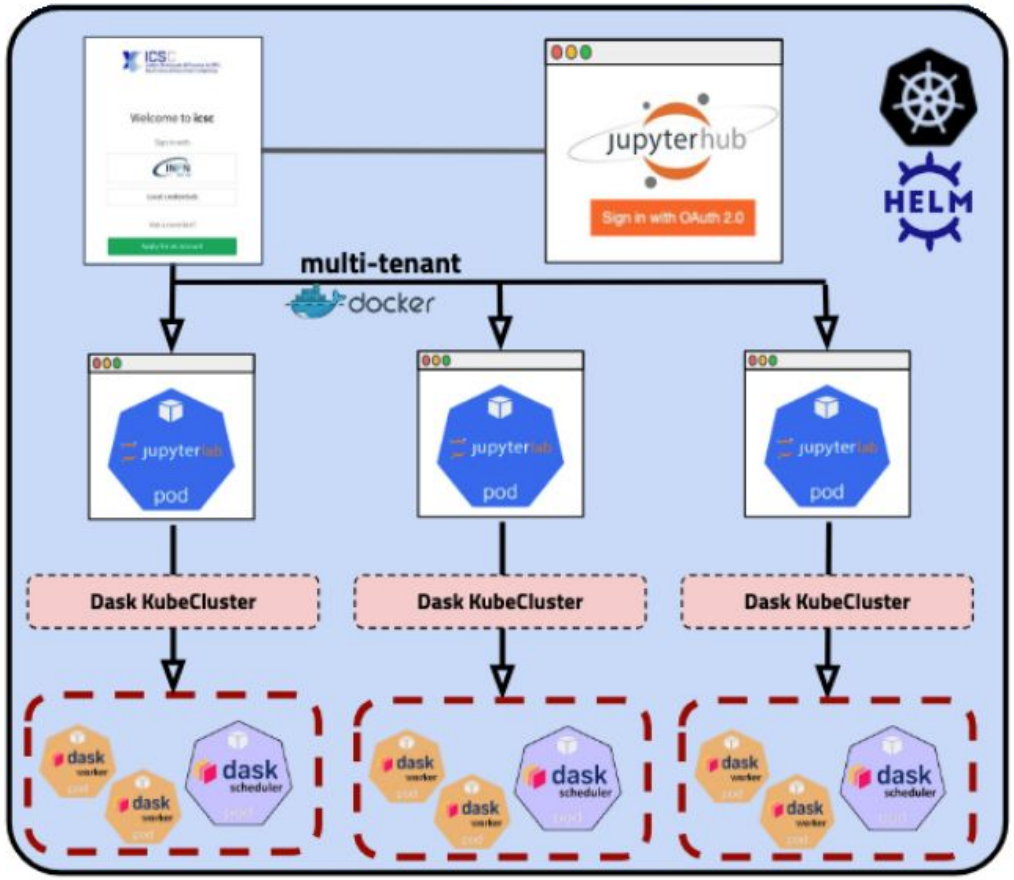

How we were until now

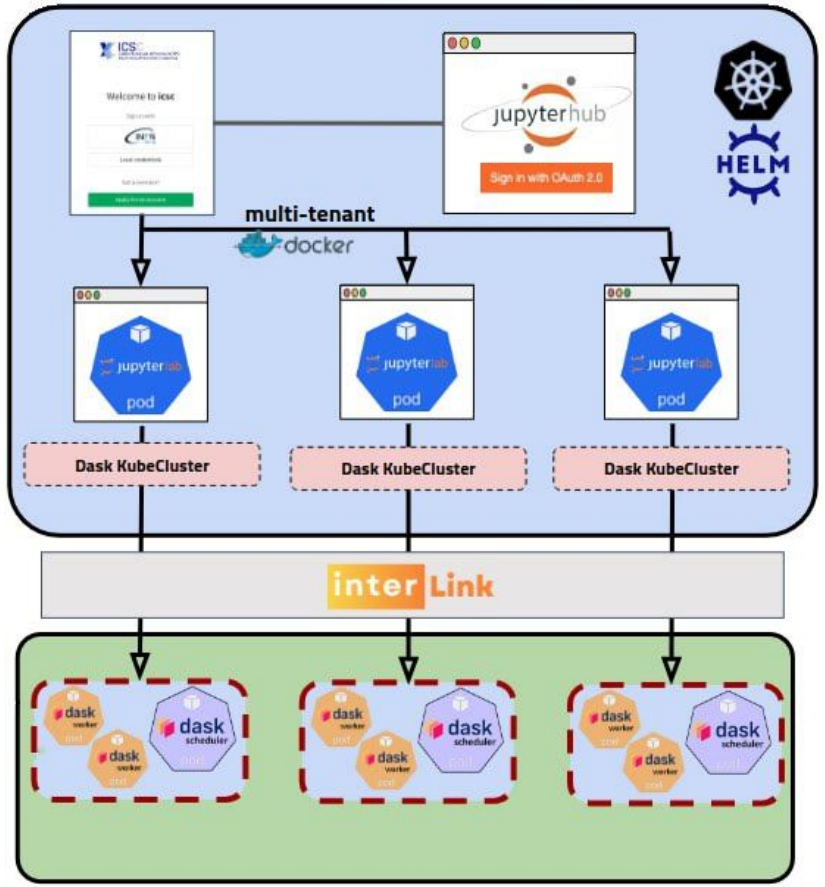

New version, also offloading to Naples resources

News from WP5

- Access is available at the usual entrypoint, with the only difference of using a **dedicated JupyterLab image** from the initial page: *ghcr.io/icsc-spoke2-repo/jlab:wp5-alma9-highrate-offload-v0.0.2-cvmfs-infn*
- Images are developed on the official ICSC-Spoke2-repo

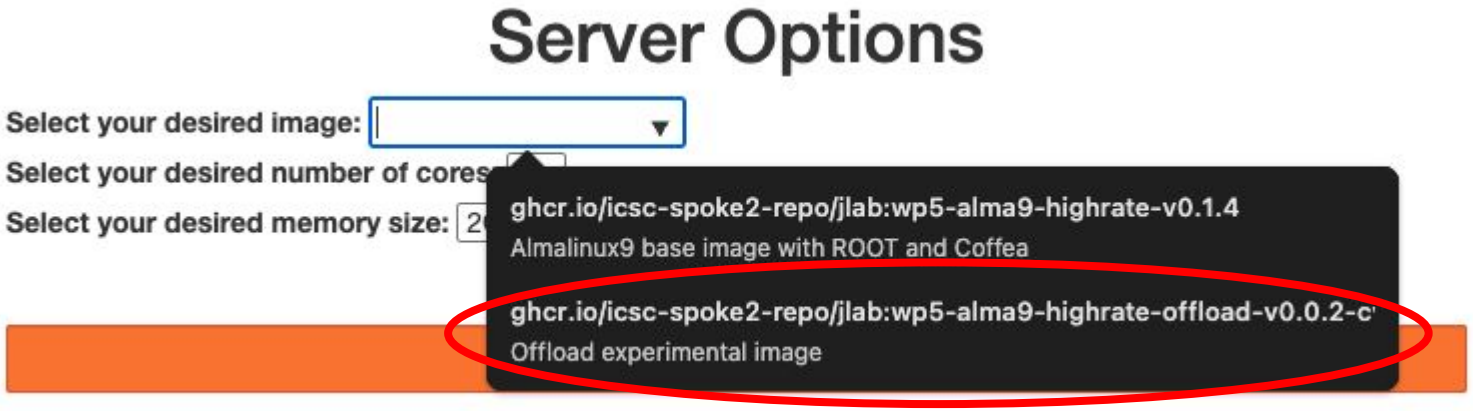

KPI - Key Performance Indicator

KPI ID	Description	Acceptance threshold	Status
KPI2.2.2.1	Implementation of *N* data analyses in the AF	$N \geq 2$	>100%
KPI2.2.2.2	Reference documentation of the AF	≥ 1 dedicated web site	100%
KPI2.2.2.3	Hands-on workshops for AF users	≥ 1 workshops	100%
KPI2.2.2.4	Scaling up the testbed AF infrastructure, serving *k* tenants, for a total of *N* data analyses	$\geq (200 \cdot N)$ cores	**>100%**
KPI2.2.2.5	Talks at conferences/workshops about AF activities	≥ 1 talk	>100%

- For the KPI2.2.2.4 definition, even the 5 initial dedicated nodes are enough to reach the planned threshold.

Work Plan for next months

- During summer, we will start benchmarking the final deployment of the platform:
  - A **first phase** (by the end of the MS10) with individual tests on the upgraded platform with the offloading on the new resources.
  - During such tests, we don't expect shared usage (i.e. queueing). Monitoring dashboard for resource consumption (CPU, memory, I/O wait, network bandwidth) will be available (on InfluxDB).
  - An operative working meeting was held last Thursday to start planning test activities and individual time windows. Documentation for access was updated accordingly.
- A **second phase** (up to end of 2025) will follow: here concurrent execution can be expected, as well as final monitoring to ensure robustness of the platform, heterogeneous benchmarking of the various analyses "flavours".

Work Plan - Phase I

- During this phase, the idea is to “book” a slot where each user can do the desired tests with no concurrent usage.
- **Coordination tool**: Use this spreadsheet to reserve the slots where you would like to test the infrastructure.
- **Users involved**: All the users that have an ongoing use-case with the platform.
- **Issue tracking**: In case of problems, fill the tracking document for logging history.

Maximum of days (per user): no more than 8-10 days.

Work Plan - Phase I

- **Metrics collection**: we should have standardised ways of collecting metrics (to avoid apple-pear situations).
  - *Timestamps*: they can be taken from the "jupyter notebook" side, with `time` or `datetime` modules. The various analyses are different in flavour, but ideally times should be taken in the same blocks:
    - Pre-dask operations (any I/O, pre-processing, or client setting);
    - dask event loop(s) (time of dask workers operation);
    - Post-dask operations (if any): plots creation, stage-out on WLCG resources.
  - *Hardware resource usage*: we have a Influx-DB instance where all the resource usage metrics are stored (CPU usage, memory, network read, Disk IO, ...)
    - This solution will be active only during Phase I of testing (to collect all the metrics);
    - The influx-DB service is not supported from WP5 officially: during the Phase I it will be provided as a best-effort solution.

Thank you

cn1-spoke2-wp2-analysisfacility@lists.infn.it (click here to subscribe or get in touch in case of issues)

ICSC Workshop on Analysis Facilities

- The workshop was held in Bologna, January 8th - 10th . Agenda
- A survey was opened (for the workshop participants) to gather feedbacks on the organisation and to collect possible new use-cases.
  - A few people replied to the survey, despite the attendance
  - Please, take some time to reply at this link (in particular, the survey part)
- We'll circulate a reminder in the upcoming days (on the mailing lists).