



WP2.2 news

Piergiulio Lenzi, Alberto Annovi
14th January 2024



Workshop on quasi interactive analysis

Held last week in Bologna:

<https://agenda.infn.it/event/44199/timetable/?view=standard>

- ⇒ Tutorials
- ⇒ Fruitful discussions
- ⇒ +25 users registered to the platform
- ⇒ +18 simultaneous connections during the hands on

Actions:

- ⇒ Scale up resources (dedicated WP5 meeting on the 17th)
- ⇒ Circulate a workshop survey



Backup

ICSC project extension

The project has been formally extended to Dic 31 2025 (+4 months).

No change to the **MILESTONE 9** schedule, which stays at the end of October

MILESTONE 10 stays at **Aug 31st 2025** BUT:

⇒ An **interim report is to be produced by mid-Milestone 10** (~February)

NEW MILESTONE 11 from 1st Sept to 31st Dec 2024

⇒ Buffer for KPIs that spill over MILESTONE 10

⇒ Single reporting document from Spoke 0

Same as last meeting



Tracking WP2.2 activities

- ⇒ Tracking of WP2.2 will happen in [this document](#). The document can be edited freely so you can update information. (Just the summary sheet is protected)
- ⇒ It contains also a list of conference opportunities
 - ✓ You are welcome to add in new conferences.

Conference	Dates	Venue	Website	comments
ANIMMA 2025	9-13/6/2025	Valencia	https://animma.com/	Deadline for abstracts: 14 Nov 2024
ACAT 2025	8-12/9/2025	Hamburg	not yet available	https://inspirehep.net/conferences/2810973



MILESTONE 9 Reporting

For the flagship PIs:

- ⇒ Expected targets for each flagship use case in [this document](#)
- ⇒ **Draft MS9 reports are available.** They are linked in next slide
- ⇒ **By Nov 12th:** reports ready to be handed to the spoke leaders.



Open Calls related to WP2

- ⇒ FaDER: Fast Data and Event Reduction: lowering data volumes in high-intensity experiments (Genova, FPGA)
 - ✓ supervisor Andrea Celentano
- ⇒ DarkSieve: Ricerca di portali a settori Dark attraverso l'identificazione di getti “boosted” originati da quark dark con tecniche di Machine Learning e nuove tecnologie ai futuri collisori adronici (Udine, UltrafastSim)
 - ✓ supervisor Antonio Sidoti
- ⇒ BOODINI: BOOsting Discoveries of New Interactions (Genova, UltrafastSim)
 - ✓ supervisor Andrea Coccaro
- ⇒ BOOST: Boosted Object and Oriented-Space Topologies from VBS@HL-LHC (Perugia, FPGA)
 - ✓ supervisor Pietro Govoni
- ⇒ SPARCh: Simulation of Photons with Accelerated Ray-tracing for Cherenkov detectors (Genova, Machine Learning)

Tracking WP2.2 activities

- ⇒ Tracking of WP2.2 will happen in [this document](#). The document can be edited freely so you can update information. (Just the summary sheet is protected)
- ⇒ It contains also a list of conference opportunities
 - ✓ **You are welcome to add in new conferences.**

Conference	Dates	Venue	Website	comments
CHEP 2024	19-25/10/2024	Cracow	https://indico.cern.ch/event/1338689/	Deadline for abstracts: 10 May
SIF	9-13/09/2024	Bologna	https://www.sif.it/attivita/congresso/110	Deadline for abstracts: 30 April



WP2 meetings

Standing call for contributions.

We no longer host flagship reports only, you can request a slot in every biweekly meeting of WP2, regardless of which flagship use case your activity is affiliated with.

Just send an email to <mailto:piergiulio.lenzi@fi.infn.it> and <mailto:alberto.annovi@pi.infn.it>

Next WP2 meeting will be **September 24th**

ICSC annual meeting

Rome, 26th-20th of September.

Agenda: <https://agenda.supercomputing-icsc.it/event/2/>

High profile meeting, mostly institutional. Only very high level scientific discussion.

⇒ One spoke-2-wide 15 min talk. Speaker Tommaso Diotallevi.



MS8 reporting - DONE

Thanks to all involved!

- ✓ TAR2.10 [MS report for UC2.2.1, including the setup of validation workflows on Cloud resources to partially automate and extend the validation campaigns. Tests of the offloading technologies developed in WP5.]
 - KPI: **Spoke2_TAR2.10.pdf** Ultrafast sim
- ✓ TAR2.11 [MS report for UC2.2.2, including a prototypal central infrastructure on available resources]
 - KPI: **Spoke2_TAR2.11.pdf** Analysis facility
- ⚠ TAR2.12 [UC2.2.3: Trigger, DAQ, on-line processing: Development of digital trigger logic for a "missing energy" experiment with a positron beam at CERN (POKER/NA64): porting of a missing momentum reconstruction algorithm to FPGA and its testing on a dedicated testbed].
 - KPI: The OPEN CALL has not started yet, so we defer the activity. Still, the Open Call has been approved and funded (see **Spoke2_TAR2.12.pdf**) FPGA
- ✓ TAR2.13 [UC2.2.3: Developing FPGA tools: Development and testing of RDMA over converged ethernet (ROCE) on FPGA for data transfer from detectors' front-end to computing servers: adapting the system developed the past year to the FPGAs being used at the moment]
 - KPI: **Spoke2_TAR2.13.pdf** FPGA
- ⚠+ ✓ TAR2.14 [UC2.2.4: Organization, in collaboration with WP4 of training events to introduce students and collaborators to heterogeneous computing and portability tools. Study of experiments' heterogeneous software stack: status and needs, updating the material harvested in this document. Defining, for each experiment, the algorithms intended to be ported within the activities of this WP. Testing the resources from INFN Cloud and Leonardo: access, nodes availability, standard workflows benchmarking].
 - KPI: the Target implied 2 different pieces of activities. The first one, about a training event, has been postponed due to
 - The non availability of experts in the MS period
 - The attempt to organise a larger event with more tutors
 - The definition of algorithms is successfully completed GPU
- ✓ TAR2.15 [MS report for UC2.2.5, including the procedures to procure and configure ARM machines in order to provide access to the experiment software and storage via a production infrastructure; selection and documentation of the workflows to be benchmarked from the most representatives; preparation of a validation strategy agreed with the experiments.].
 - KPI: **Spoke2_TAR2.15.pdf** ARM

WP5 workshop on data management

If you haven't attended, please check the slides and recording at:

<https://agenda.infn.it/event/42119/timetable/?view=standard>



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



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

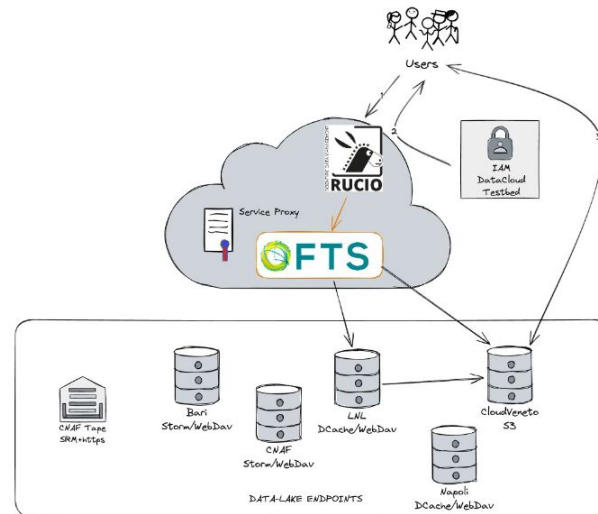
The testbed in DataCloud

We have chosen to start the experimentation by integrating **de-facto** standard tools already in use in scientific realities close to us and that we are familiar with (e.g., LHC):

- **Rucio+FTS** (Data manager)
- **IAM** (AuthN/Z)
- Metadata Catalog: embedded in **Rucio**

6 heterogeneous storage systems of INFN:

- Qos (disk, tape);
- One storage with S3 protocol on ceph @CloudVeneto
- Three storages with WebDav protocol
 - Two based on STORM (CNAF, Bari)
 - Two on dCache (LNL, Naples)
- One tape endpoint @CNAF



Federare lo storage distribuito nazionale, D. Ciangottini, 2023

Open Calls status

N. area	Area tematica	Vincoli	PROPOSTE
PUB1	Development of data acquisition and trigger strategies in the context high background, rare event experiment, using acceleration on hybrid CPU/FPGA architectures, in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	1
PUB2	Development of a multipurpose/multi experiment high level cloud infrastructure for small/medium astroparticle experiments, on the infrastructure of ICSC and for the benefit of Spoke 2 use cases	Per beneficiari pubblici	
PUB3	Development of a GPU library for massively parallelized simulations of QCD and QCD+QED on the lattice in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	1
PUB4	Development and performance optimization of typical simulation codes in HEP and Astro, for the use cases of Spoke 2	Per beneficiari pubblici	4
PUB5	Advanced algorithms for GW experiments (Virgo and ET) in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	2
PUB6	Studying the quantum Boltzmann equation neutrino flavour conversion in a dense environment in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	
PUB7	Development of algorithms for boosted topologies at LHC/FCC, in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	3
PUB8	Accelerated analysis of Astrophysical data in the Photon Counting Big Data era, in the context of the activities of the Spoke2 in ICSC	Per beneficiari pubblici	

N. area	Area tematica	Vincoli	PROPOSTE
IND1	Data Management PoCs/pilots with Spoke 2 solutions in industrial environments (datalake, distributed computing, ...)	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	3
IND2	Solutions, implementation and deployment of a computing PoC based on the ARM architecture	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	0
IND3	Porting and optimization (on GPU, on FPGA, on CPU) of algorithms of interest of AstroParticle experiments	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	
IND4	Realization of tools and pilots for the Space Economy domain, using the infrastructure of the ICSC (portals, market exchange, algorithms, services, ...)	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	
IND5	Heterogeneous industrial use cases on the Spoke 2 platforms (GPU + FPGA + ARM)	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	
IND6	Profiling, code engineering and code quality on Spoke 2 code repositories	Per beneficiari privati; non utilizzabile "Ricerca Fondamentale"	



Resource Allocation Board

- ⇒ All requests collected were approved.
- ⇒ Great success, thanks to the PIs for their well documented requests.
- ⇒ Flagship use cases PI have been in contact with CINECA and INFN to gain access to the resources, more in the talks
 - ✓ CINECA resources might come first
- ⇒ **WARNING:** allocation of resources in future requests are likely going to be more selective

Per referenza

	Durata temporale	Mese di progetto	Scadenze per invio report scientifico ad HUB
Milestone 6	Mag 23 – Ago 23	M9-M12	DONE!
Milestone 7	Set 23 – Feb 24	M13-M18	20 Marzo 2024
Milestone 8	Mar 24 – Giu 24	M19-M22	15 Luglio 2024
Milestone 9	Lug 24 – Ott 24	M23-M26	20 Novembre 2024
Milestone 10	Nov 24 – Ago 25	M27-M36	20 Marzo 2025 (report intermedio) Report finale a Settembre 2025 (deadline tbc)

Reporting

New procedure agreed with the referees:

For each project milestone we will need to produce a report for each flagship UC.

Milestones are:

- ⇒ Month 18 (Feb '24) → Milestone 7
- ⇒ Month 22 (June '24) → Milestone 8
- ⇒ Month 26 (Oct '24) → Milestone 9
- ⇒ Month 36 (Aug '25) → Milestone 10

Every two months an interim short report is needed as well for every flagship UC

Flagship use cases

Flagship documents available at these links:

- ⇒ **Quasi interactive analysis of big data with high throughput**
 - ✓ Tommaso Diotallevi (UniBo), Francesco Gravili (UniSalento)
- ⇒ **Advanced ML: flash simulation and other bleeding edge applications**
 - ✓ Lucio Anderlini (INFN Fi)
- ⇒ **Development of ultra-fast algorithms running on FPGAs**
 - ✓ Bernardino Spisso (UniNa), Simone Gennai (INFN MiB)
- ⇒ **Porting of algorithms to GPUs**
 - ✓ Adriano Di Florio (Poliba)
- ⇒ **Physics validation of reconstruction code on ARM**
 - ✓ Francesco Noferini (INFN Bo)

Thanks to the work of the PIs over the summer and to all of your contributions



Flagship UC mailing lists

- ⇒ Quasi interactive analysis of big data with high throughput
 - ✓ [\[cn1-spoke2-wp2-analysisfacility\]](#), [subscribe](#)
- ⇒ Advanced ML: flash simulation and other bleeding edge applications
 - ✓ [\[cn1-spoke2-wp2-flashsim\]](#), [subscribe](#)
- ⇒ Development of ultrafast algorithms running on FPGAs
 - ✓ [\[cn1-spoke2-wp2-fpga\]](#), [subscribe](#)

Please subscribe to the list that is relevant for the work you are doing in ICSC.

ICSC spoke 2 github organization

Available at: <https://github.com/ICSC-Spoke2-repo>

People are encouraged to add repositories with software developed in the context of ICSC - spoke 2. Send us a request for the creation of repos/moving of repos

Important: if you are not the owner of a repository (e.g. because it is the repository of an experiment and it contains code other than the one developed in the context of ICSC), we'd still be interested in forking the repo in this organization

Having spoke 2 code in this organization is important for reporting/auditing from the referees