From AMS to HERD: status and exploration of CLOUD solutions for cosmic ray data analysis

V. Formato

Workshop CCR – Paestum - 23/05/2022

INFN ROMA TOR VERGATA

Outline

- The AMS and HERD experiments and their computing requirements
- The challenges we face in guaranteeing fair and seamless access to the experiment computing resources / web infrastructure to a whole collaboration
- The impact on / role of INFN-Cloud in solving these challenges
- Exploration of a per-experiment based IT infrastructure

Who (acknowledgements)

- Perugia:
  - D. Spiga
  - M. Tracolli
  - D. Ciangottini
  - M. Mariotti
- CNAF:
  - D. Cesini
  - A. Ceccanti (👋🥲)
  - M. Vilaça Pinheiro Soares
  - L. Morganti
  - S. Stalio
  - C. Duma
  - …
- AMS/HERD team:
  - V. Formato (RM2)
  - M. Duranti (PG)
  - N. Mori (FI)
- Bari:
  - G. Donvito
  - M. Antonacci
  - S. Nicotri
  - …

(Many of the people that helped are also on the INFN-Cloud team)

History

EOSC-HUB / DODAS tests (2017): AMS-02 as a pilot use-case for cosmic-ray experiments

ASI farm integration: Test deployment of HTCondor on demand on dedicated resources, demo of flocking between two separate batch systems

Early adoption of INFN-Cloud: HERD among the first teams to bet on new tools and infrastructure

DICE project users: Testing storage and data-access solutions

Since early 2017 we have been collaborating in efforts and projects that later became part of the INFN-Cloud infrastructure.

AMS-02

- Installed on the International Space Station, ISS, on May 19, 2011
- Operating 24h/day, 365d/year, since the installation
- 300k readout channels + 1500 temperature sensors
- Acquisition rate up to 2kHz
- More than 600 microprocessors to reduce the rate from 7 Gb/s to 10 Mb/s
- 4 Science Runs (DAQ start/stop + calibration) per orbit: 1 Science Run = ~ 23 minutes of data taking
- On May 2019, ~135 billion triggers acquired
- 35 TB/year of raw data

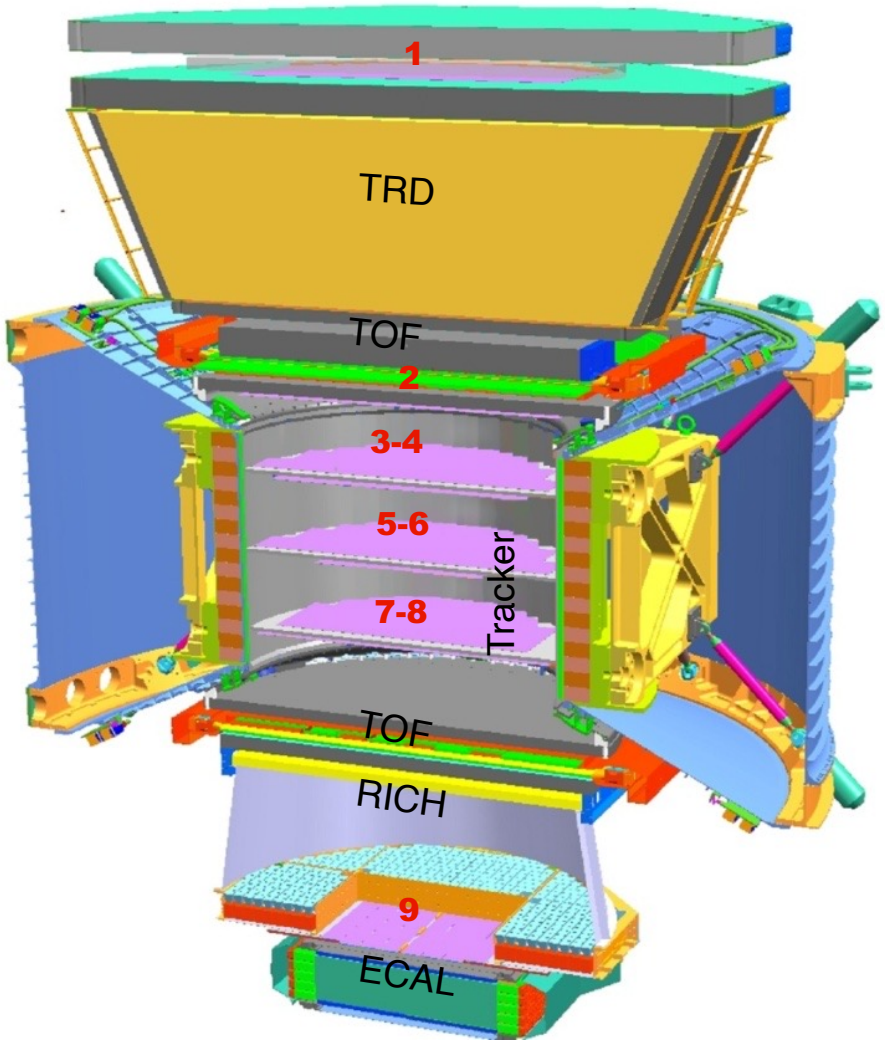

AMS-02

In the last years we learned a lot leveraging the "stable" model of AMS-02 to explore new tools for resource exploitation (in a good way...)

Now we want to consolidate all the lessons learned for the sake of future experiments

(and in addition, finally put in production our resources in ASI)

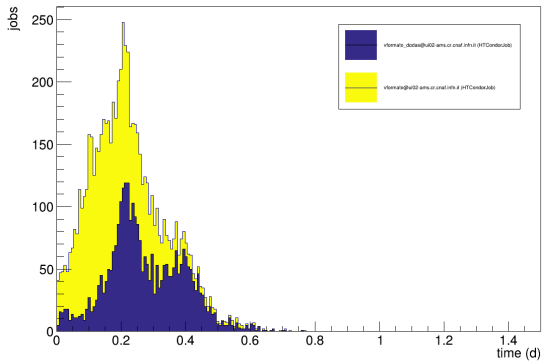

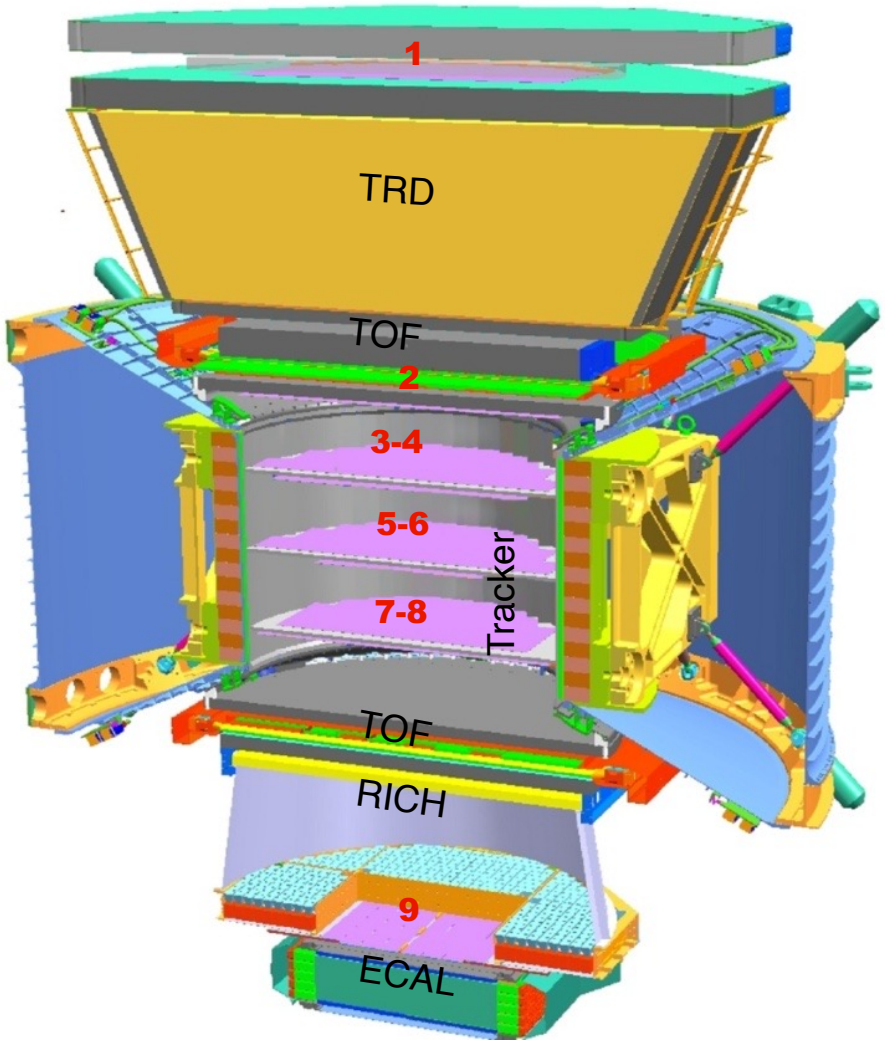

HERD

- Flagship and landmark scientific experiment, China-led large international collaboration
- Main Scientific Objectives:
  - Dark matter: dark matter search with unprecedented sensitivity
  - Cosmic-ray: Precise cosmic ray spectrum and composition measurements up to the knee energy
  - Gamma-ray: Gamma-ray monitoring and full sky survey
- Foreseen to operate in space, on board the Chinese Space Station
- Charged CR physics but also γ-ray physics
- ~ O(1M) read-out channels
- The detector is designed to be "isotropic" and accept CR from all (5) the sides

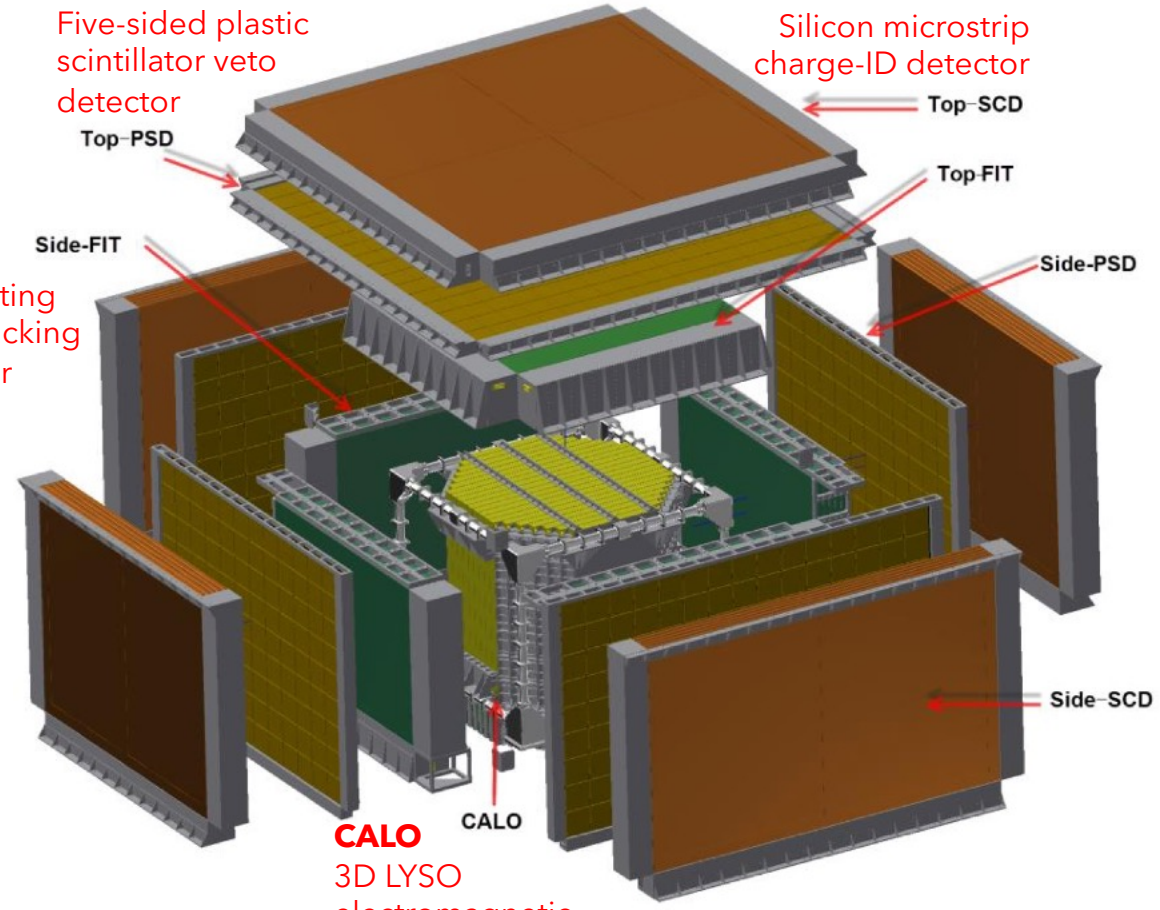

HERD – computing overview

We are currently designing the HERD computing model from scratch

- Pros:
  - There is no pre-existing structure / architecture. We are free to design the model as we see fit
- Cons:
  - We (Nicola, Matteo, Valerio) never designed a computing model before
  - Exploring new solutions always comes with a overhead and *trial and error* is a natural part of the process
  - Leveraging the expertise of our wise colleagues 🙂

* The analysis workflow is not that different from AMS, except it relies only on simulated data. Most of the work is done directly by the end-users, running simulations at CNAF and analyzing the results

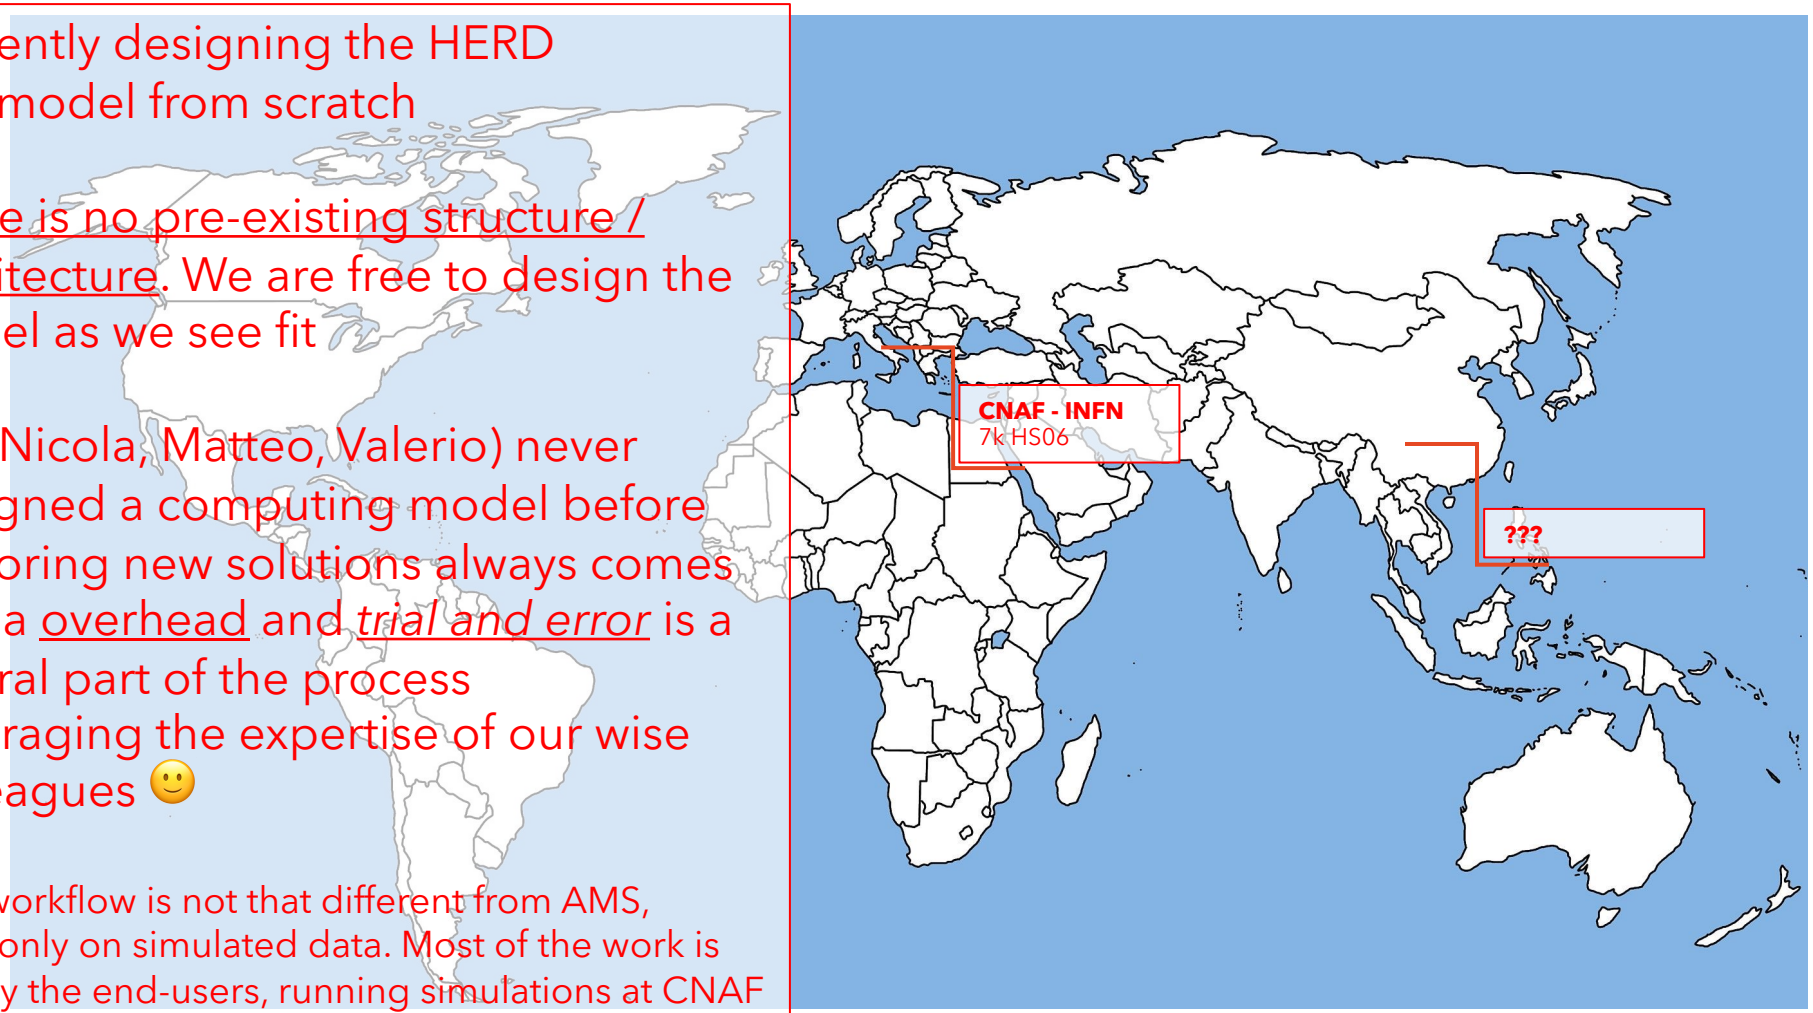

HERD – computing overview

For all these reasons INFN-Cloud seems the natural choice for prototyping all the bits and pieces that will become the HERD computing model

We are confident that all the tests and experimentation done in the AMS-02 context, and that sooner or later became part of the INFN-Cloud suite, will allow us to reach our goals for HERD.

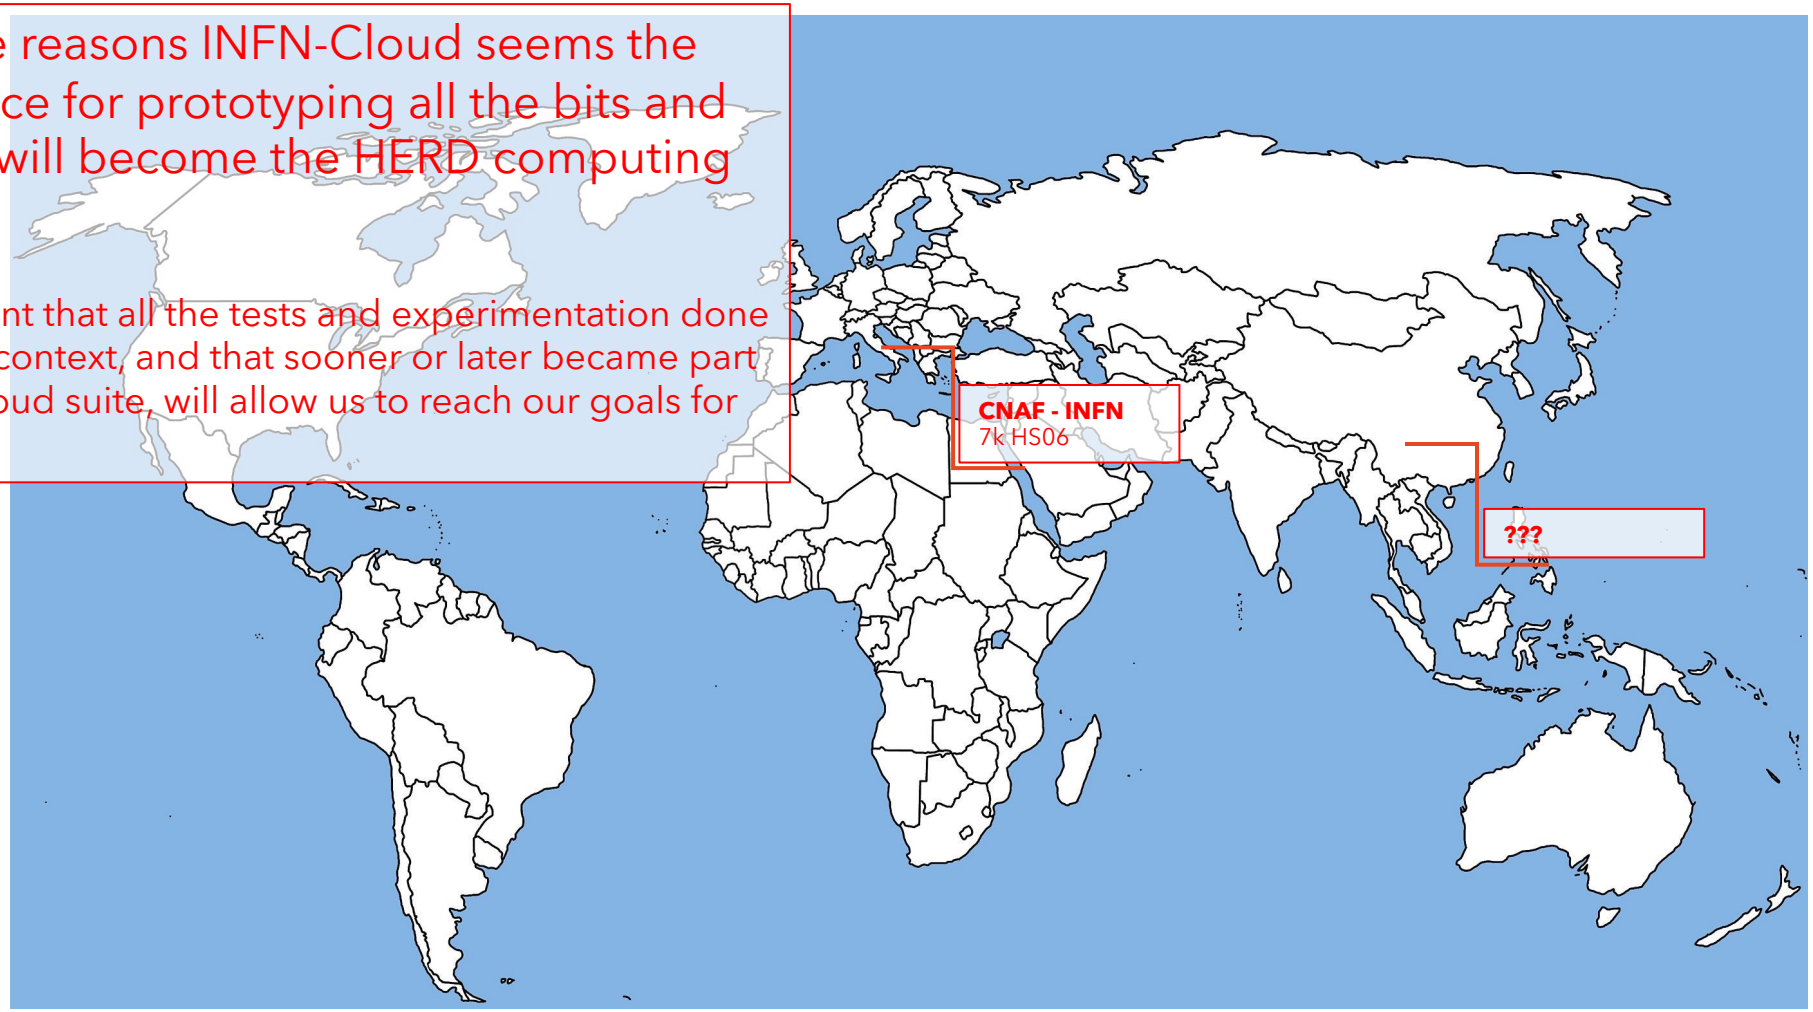

HERD – resource access

First choice: we want everyone in the collaboration to have access to the experiment resources.

This boils down to requiring a centralized authentication service, defining who is a member of the collaboration and who is not.

Issues:

- Not everyone has a CERN account
- Not everyone has a INFN-AAI account
- Not everyone has access to Google

Solution: we deploy a dedicated IAM instance

To ease login from different institutions and to comply with security rules, we feature a eduGAIN login option.

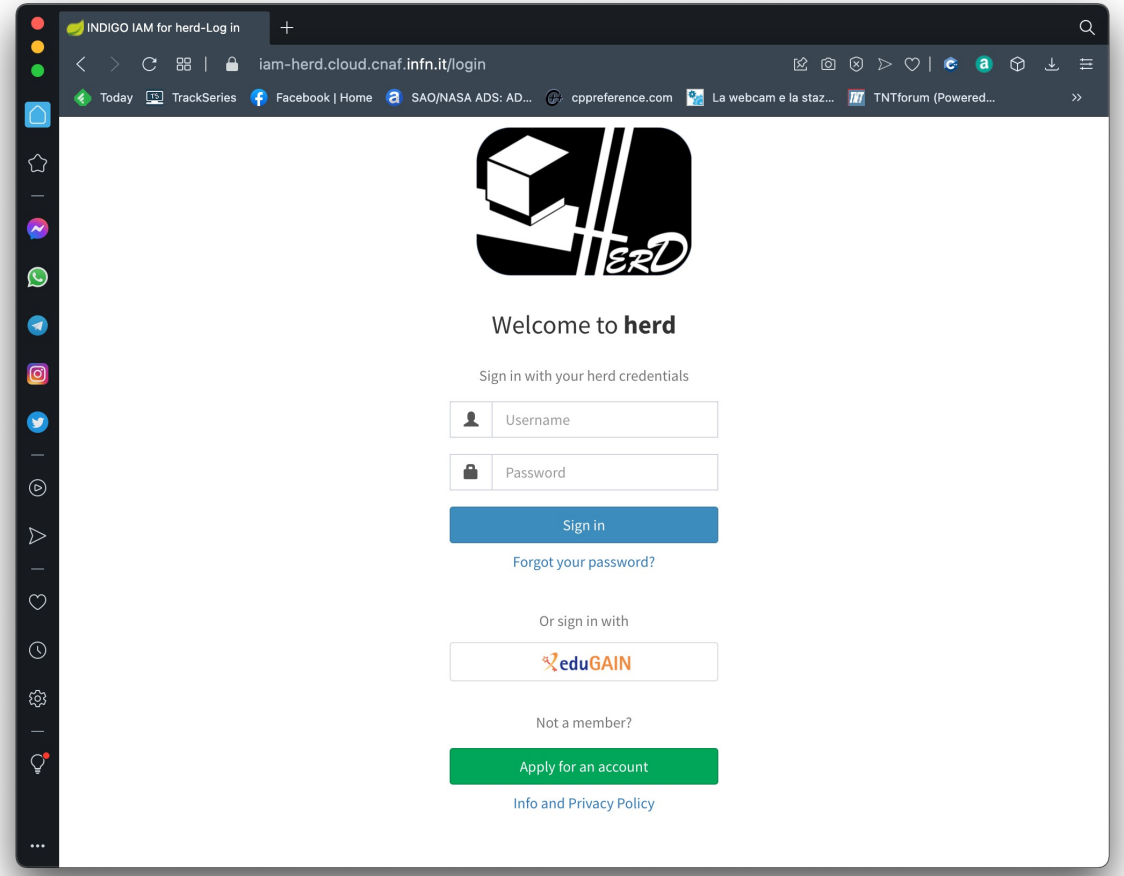

HERD – batch system

Work is currently done in a traditional way:

First-time users request a CNAF account, then login and familiarize with the batch system and the local installation of the HERD software.

Our software stack has been manually installed and mantained on our gpfs area with a rough CD based on gitlab-runner.

We are currently testing the HTCondor on-demand batch system provided by INFN-Cloud, using our experiment IAM instance for authentication.

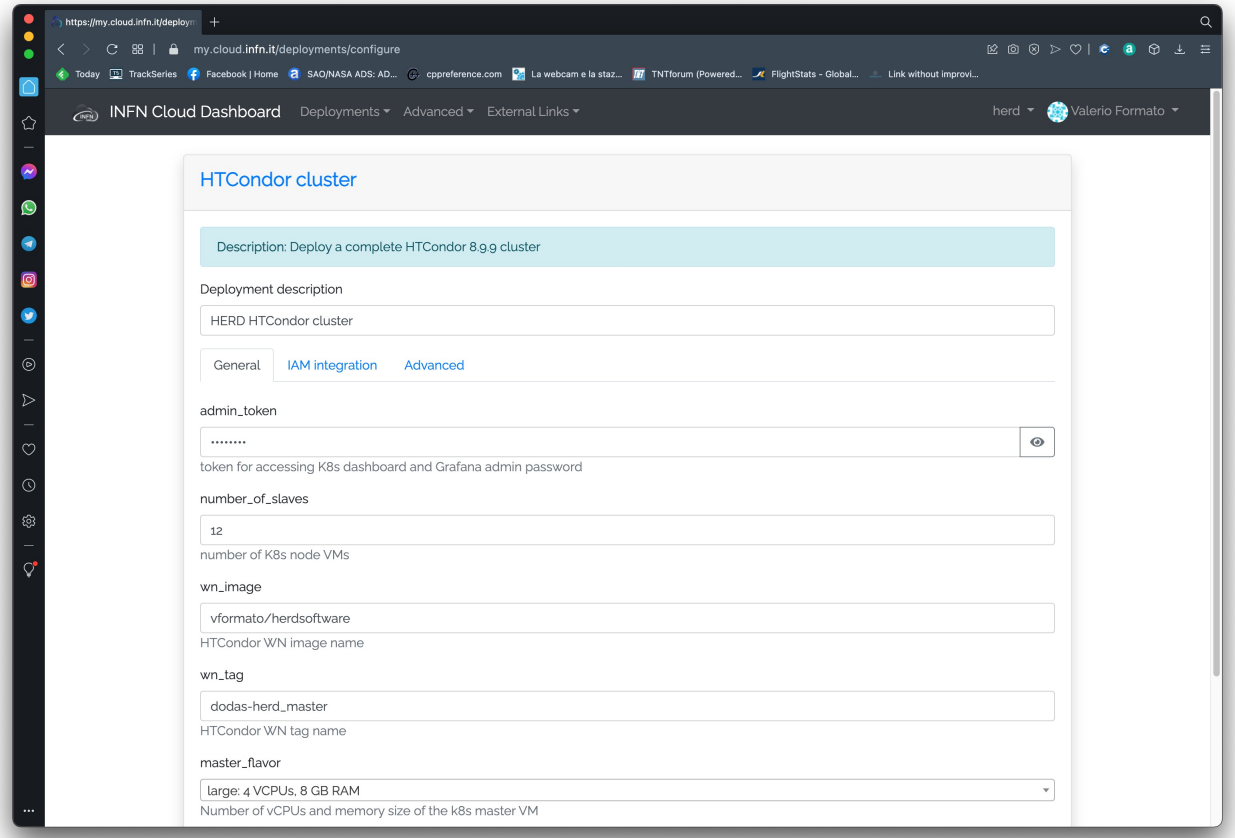

HERD – batch system

Work is currently done in a traditional way:

First-time users request a CNAF account, then login and familiarize with the batch system and the local installation of the HERD software.

Our software stack has been manually installed and mantained on our gpfs area with a rough CD based on gitlab-runner.

We are currently testing the HTCondor on-demand batch system provided by INFN-Cloud, using our experiment IAM instance for authentication.

We have a HERD tenant with some allocated resources, where we deployed a HTCondor cluster.

We were able to reproduce the whole analysis workflow, and we are currently evaluating how to integrate cloud storage into it.

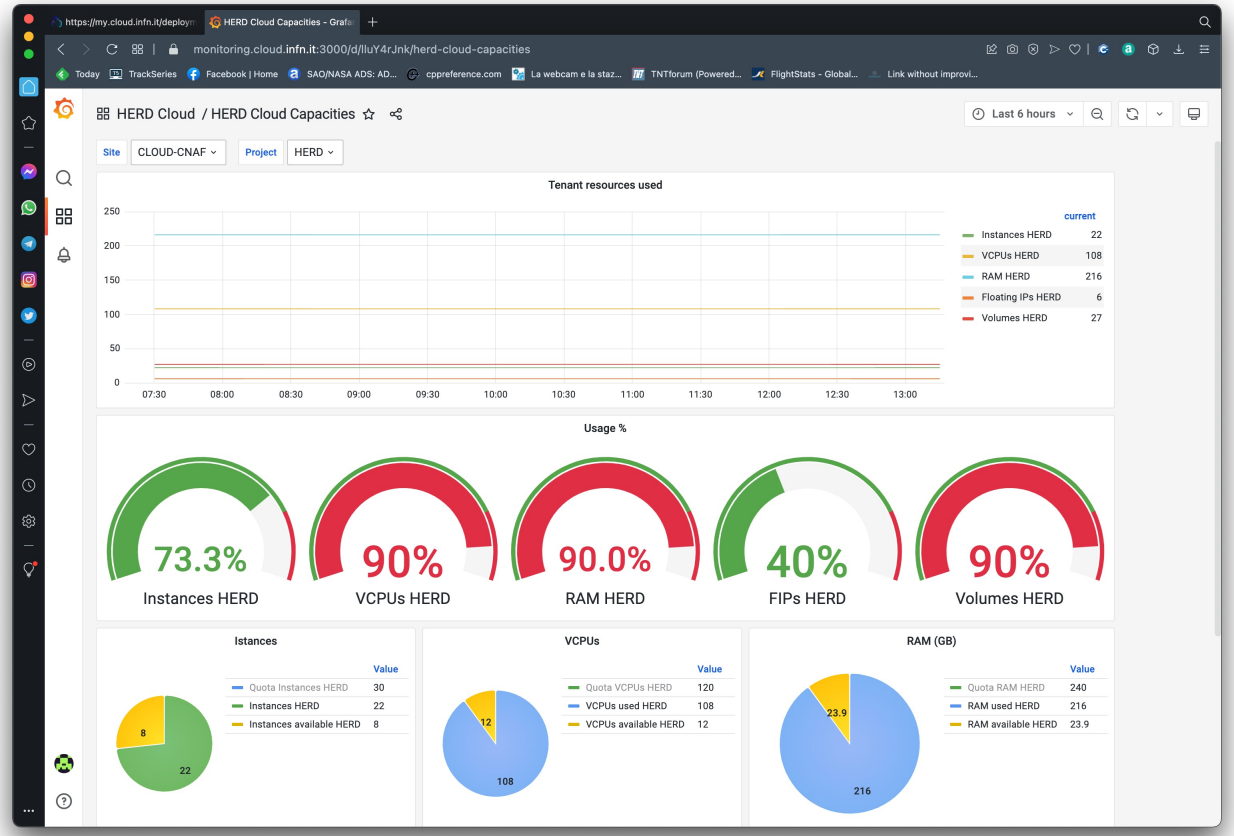

HERD – batch system

Work is currently done in a traditional way:

First-time users request a CNAF account, then login and familiarize with the batch system and the local installation of the HERD software.

Our software stack has been manually installed and mantained on our gpfs area with a rough CD based on gitlab-runner.

We are currently testing the HTCondor on-demand batch system provided by INFN-Cloud, using our experiment IAM instance for authentication.

We have a HERD tenant with some allocated resources, where we deployed a HTCondor cluster.

We also want to try integrating this with the INFN farm in ASI (where another instance will run soon) and potentially with our resources at CNAF T1.

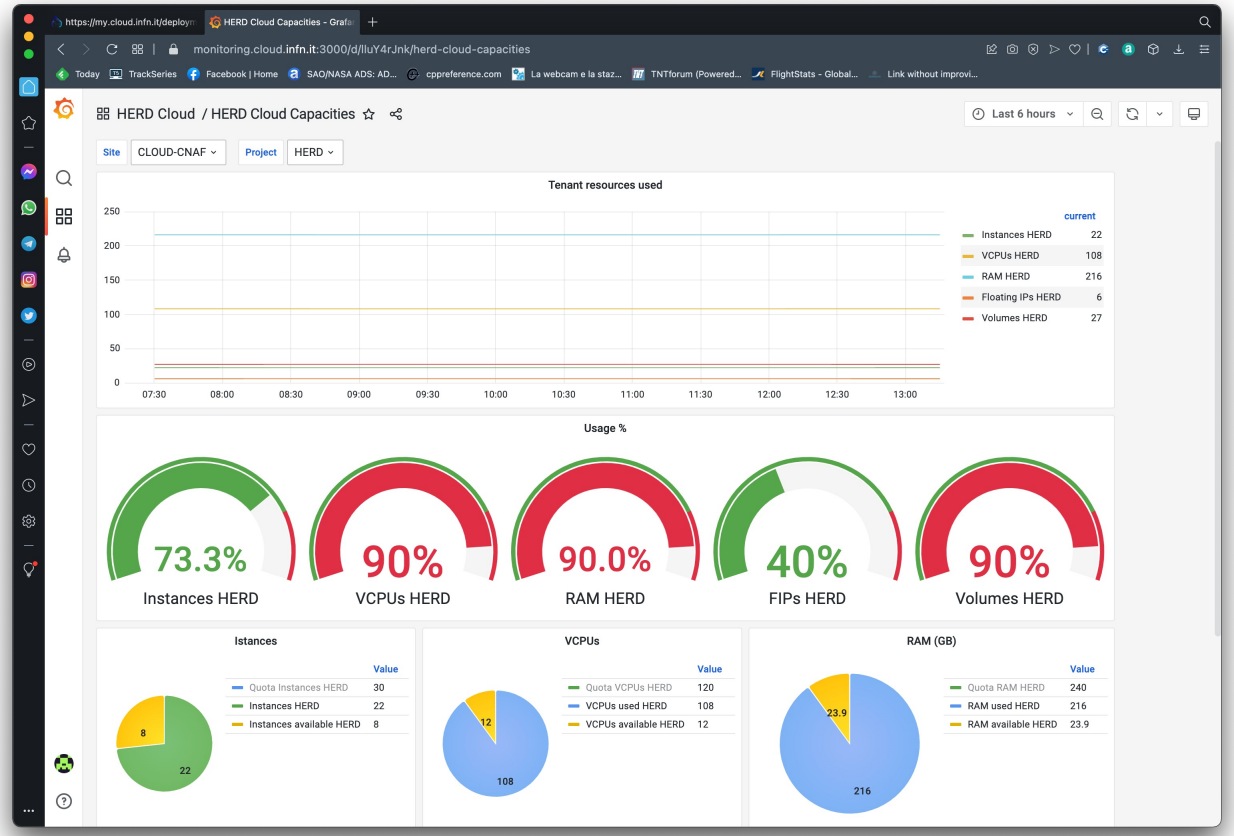

HERD – services

Requirement: we need several services to ease scheduling of day-to-day operations and/or meetings, activities, as well as document sharing, and more.

Issue: Almost every mainstream tool is unavailable to colleagues in China.

Solution: self-hosting of all services:

- Calendar (based on Nextcloud)

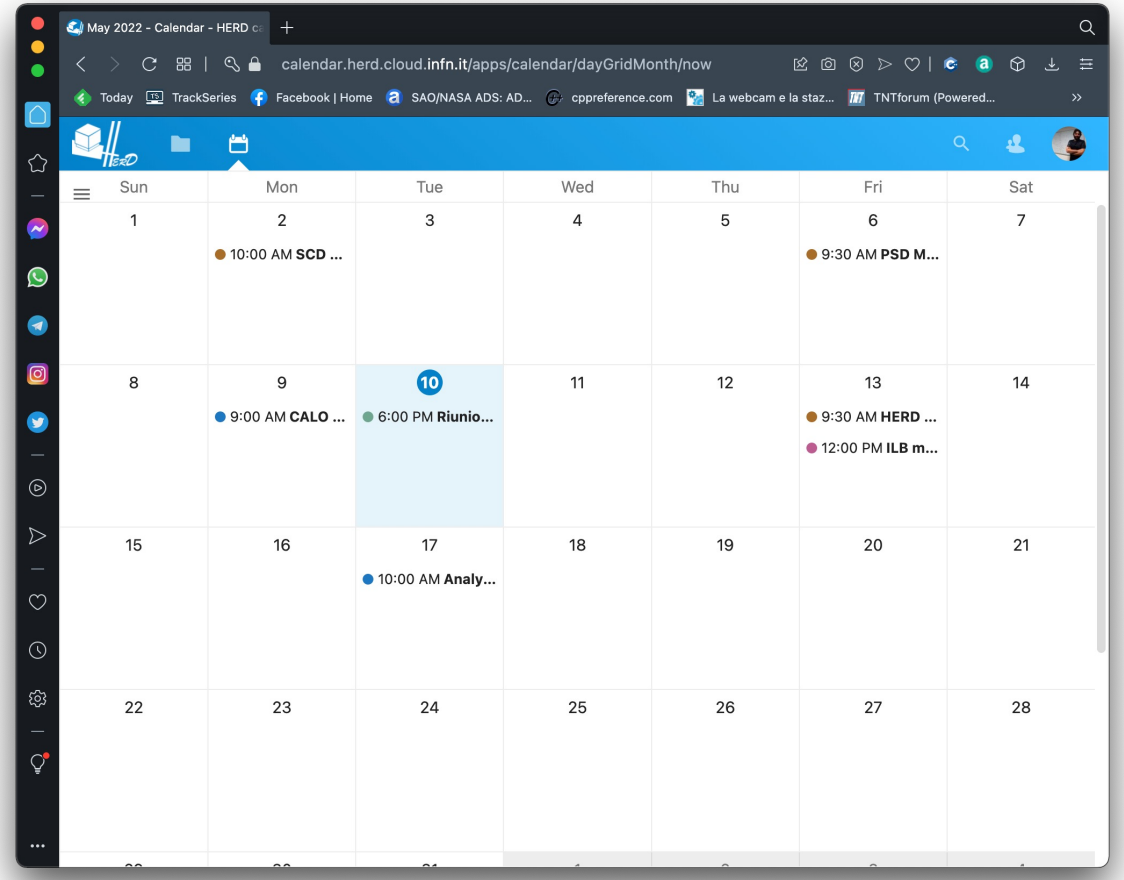

HERD – services

Requirement: we need several services to ease scheduling of day-to-day operations and/or meetings, activities, as well as document sharing, and more.

Issue: Almost every mainstream tool is unavailable to colleagues in China.

Solution: self-hosting of all services:

- Calendar (based on Nextcloud)
- Document server (based on Nextcloud)

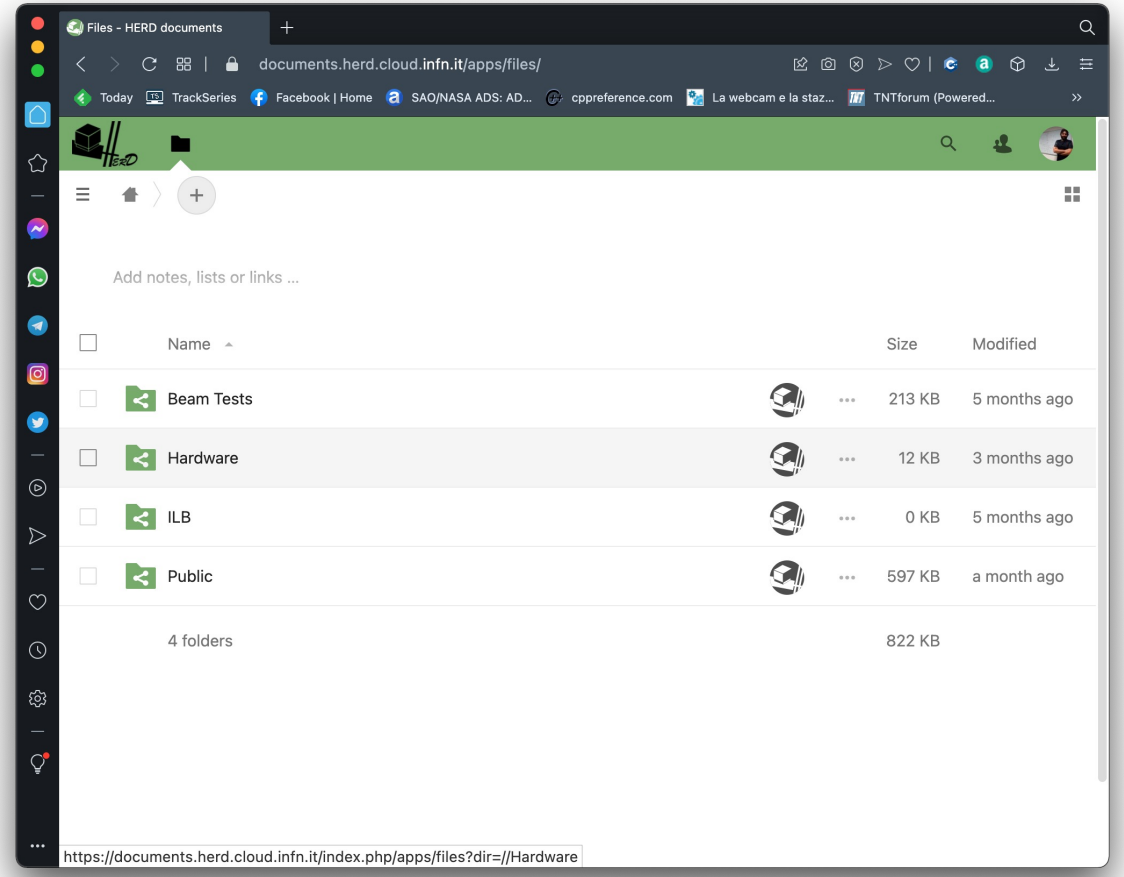

HERD – services

Requirement: we need several services to ease scheduling of day-to-day operations and/or meetings, activities, as well as document sharing, and more.

Issue: Almost every mainstream tool is unavailable to colleagues in China.

Solution: self-hosting of all services:

- Calendar (based on Nextcloud)
- Document server (based on Nextcloud)
- Experiment website (with dedicated restricted area for internal documentation, behind IAM login)

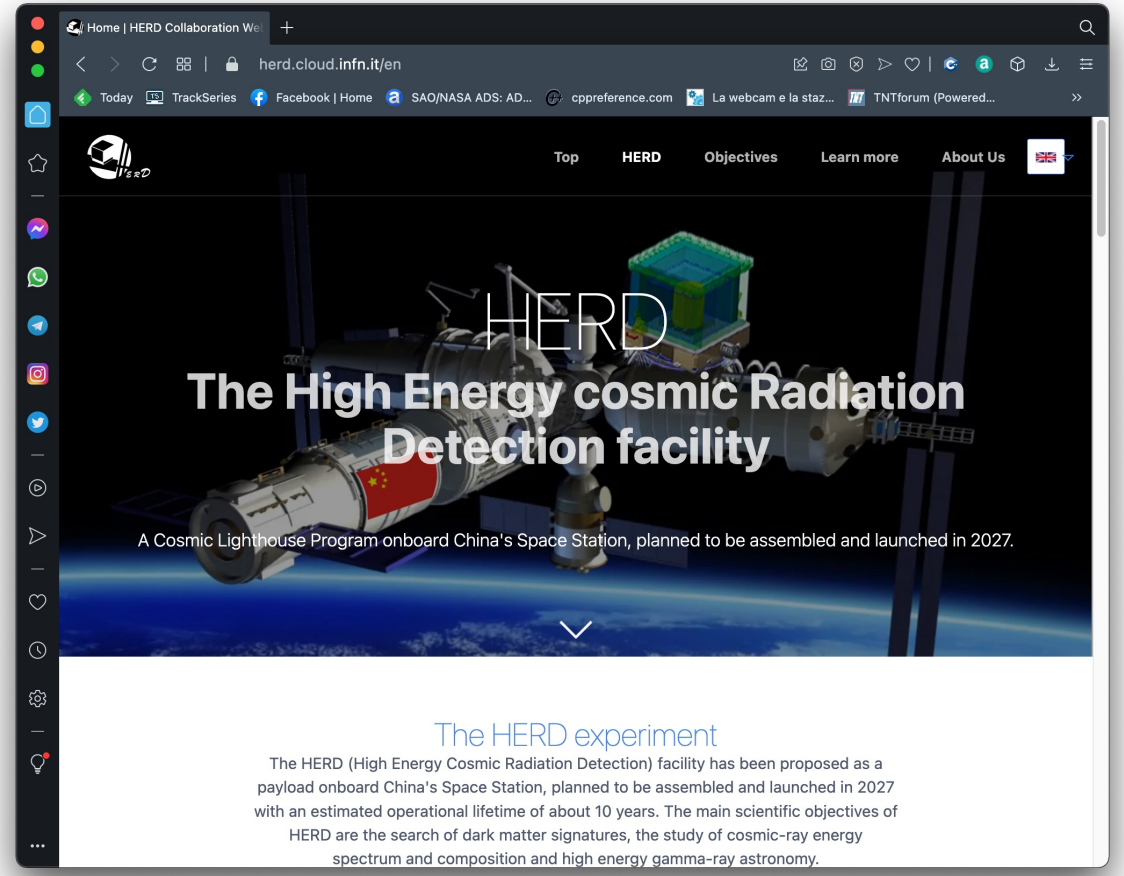

HERD – services

Requirement: we need several services to ease scheduling of day-to-day operations and/or meetings, activities, as well as document sharing, and more.

Issue: Almost every mainstream tool is unavailable to colleagues in China.

Solution: self-hosting of all services:

- Calendar (based on Nextcloud)
- Document server (based on Nextcloud)
- Experiment website (with dedicated restricted area for internal documentation, behind IAM login)

Currently testing also:

- Experiment-wide chat service (based on Matrix)

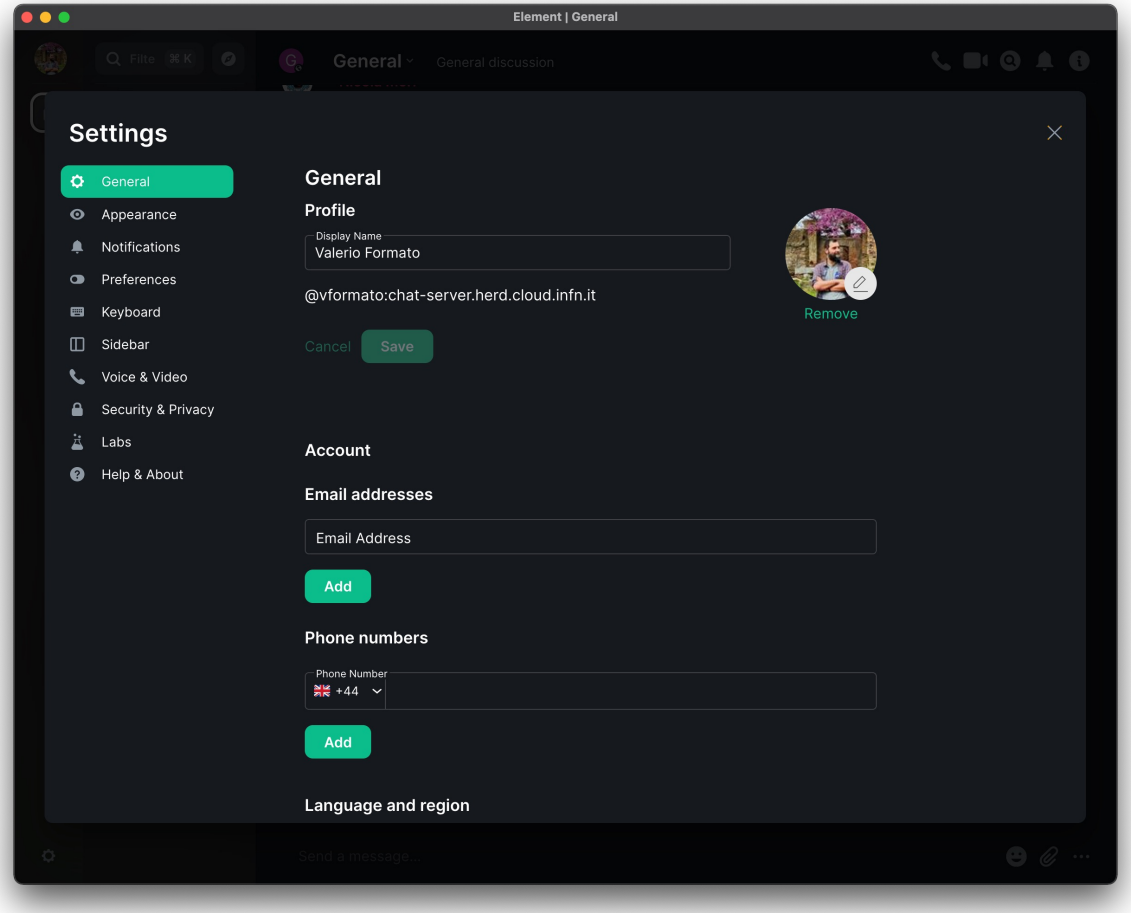

HERD – services

Requirement: we need several services to ease scheduling of day-to-day operations and/or meetings, activities, as well as document sharing, and more.

Issue: Almost every mainstream tool is unavailable to colleagues in China.

Solution: self-hosting of all services:

- Calendar (based on Nextcloud)
- Document server (based on Nextcloud)
- Experiment website (with dedicated restricted area for internal documentation, behind IAM login)

Currently testing also:

- Experiment-wide chat service (based on Matrix)
- Gitlab instance (with runners)

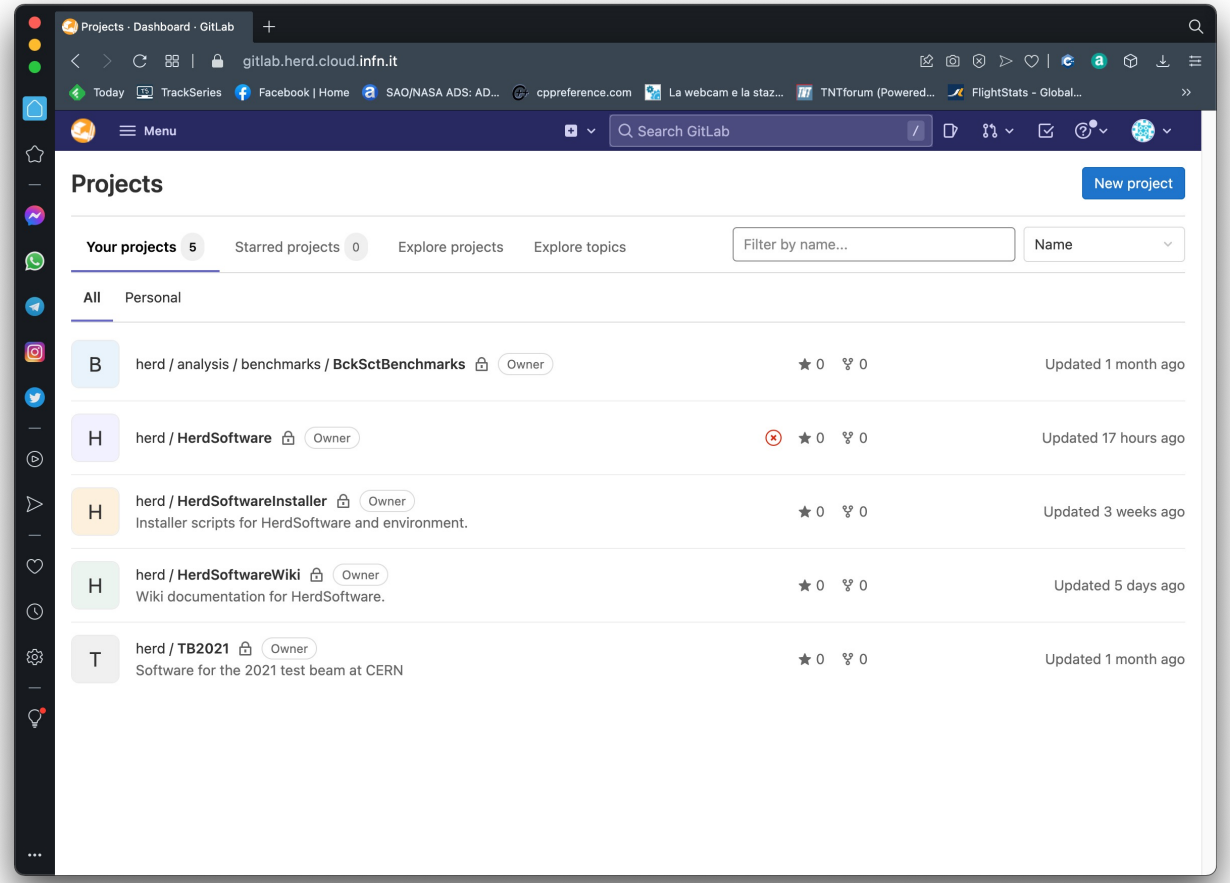

HERD – storage

Storage solution: For testing purposes our storage is provided by a Ceph cluster hosted at CNAF with a Minio gateway for access via S3.

Storage access is also behind IAM authentication

This storage is used for:

- Backup of service data (databases, artifacts, ...)
- Persistence for document service
- Storage for HTCondor jobs

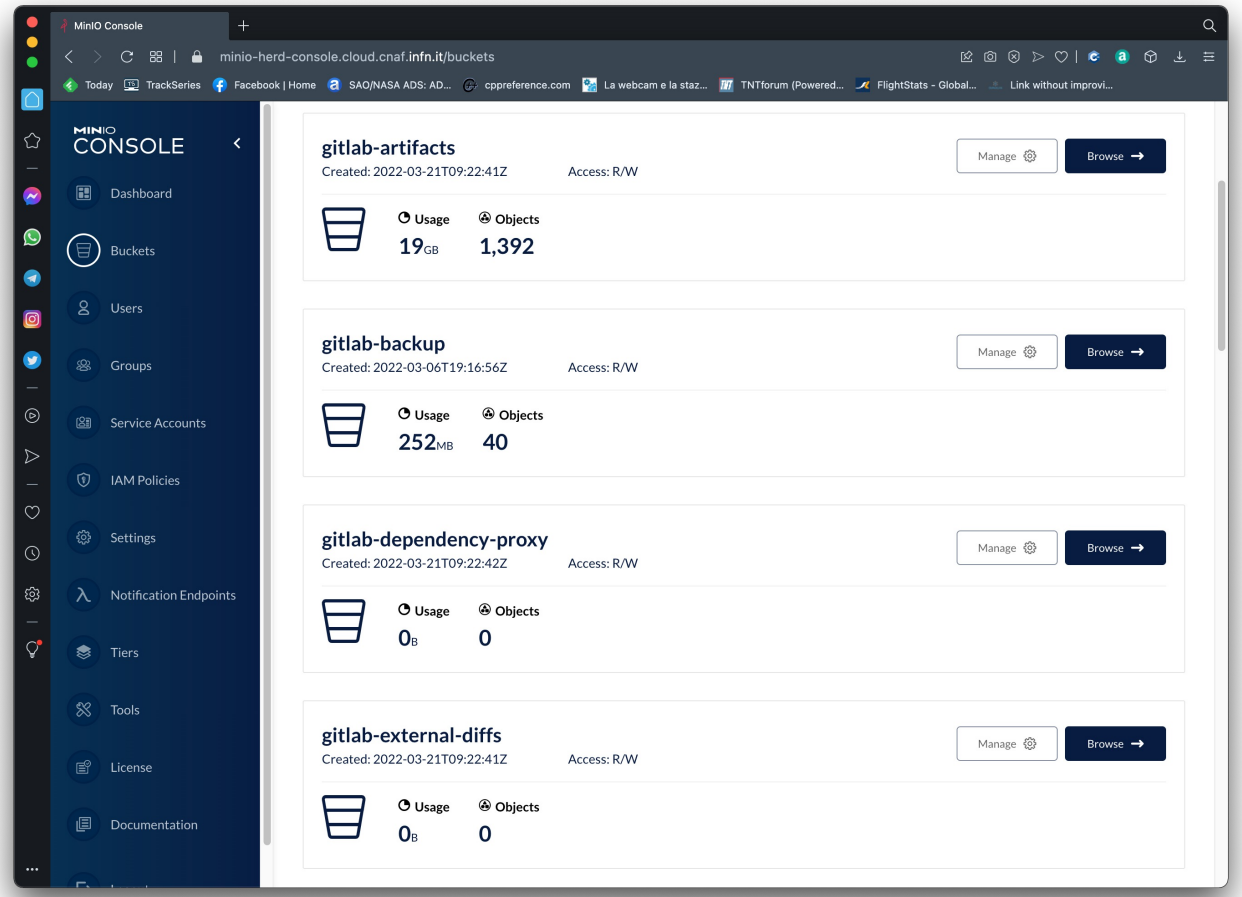

HERD – infrastructure model

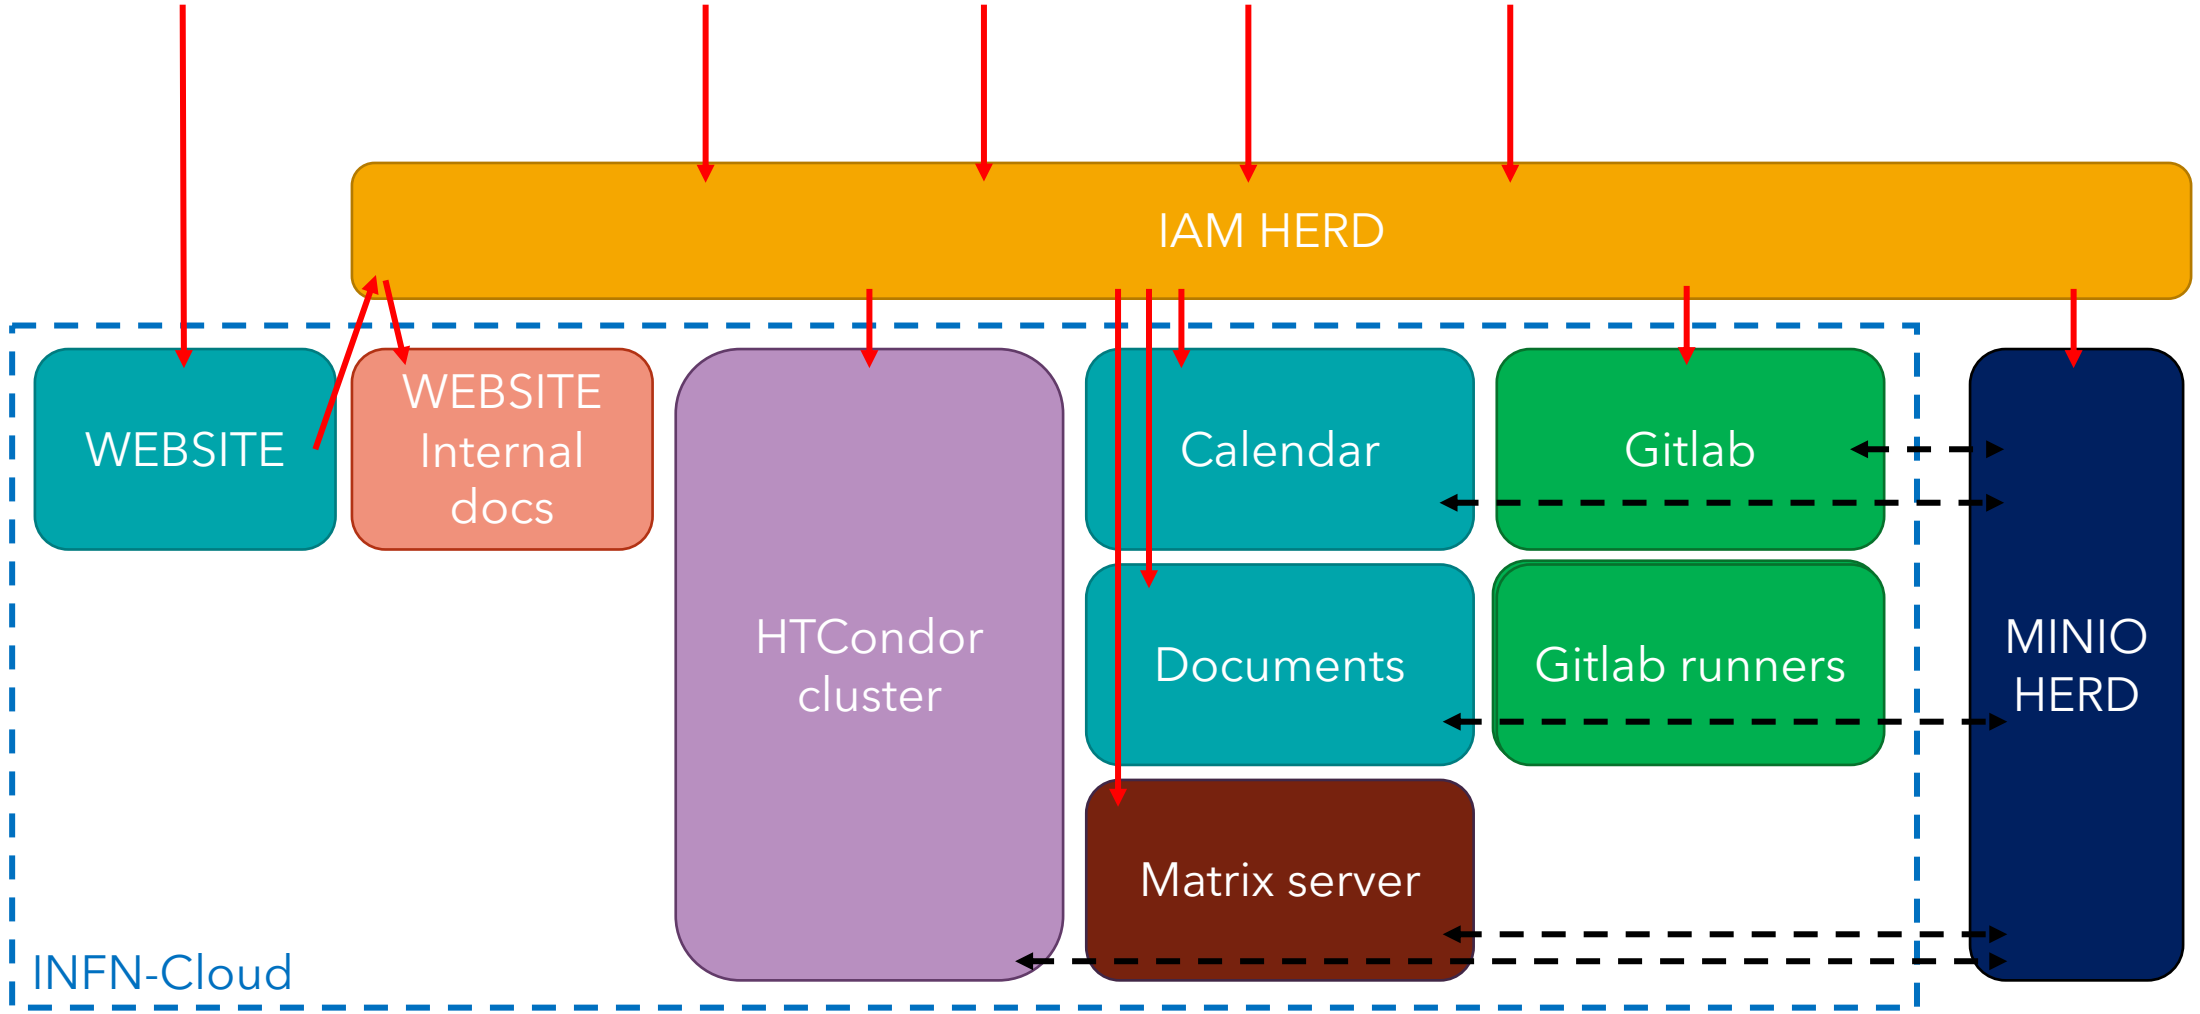

Lessons learned

- The cloud paradygm is enabling very small teams to build up a comprehensive set of experiment resources and to properly plan a computing model for small experiments.
- Some predisposition is required from the experiment personnel, however the learning curve for all these new tools is not as steep as it can look initially.
- However, when dealing with complex or delicate services (i.e. authentication, storage, ...) the help and support from experts is of paramount importance.

Feedback and next steps

- Needed:
  - CVMFS as-a-service.
    To continue testing the HTCondor service (and eventually use it in production) we need a CVMFS server for software distribution.
- Desirable:
  - Self-managed DNS on select subdomains
    To ease the workflow of deploying new services or maintaining existing ones
- Feedback:
  - Support team with high availability for critical, non-self-managed parts of the infrastructure (*e.g.* storage)
    We had several issues with our Minio backend which we couldn't debug/fix on our own. We relied on a single contact person during the test phase but this solution clearly doesn't scale to production.
  - The whole INFN-Cloud infrastructure feels a bit "user-centric" rather than "team-centric"
    For example: within the HERD tenant each user can see and manage only his own deployments
- Next:
  - Continue testing and improve the HTCondor workflow for analysis
    Now testing S3 storage solutions, need CVMFS integration, flocking, user-mapping, ...
  - Backport hand-made deployments in INFN-Cloud dashboard (*e.g.* Gitlab as-a-service)

Conclusions

- We (*i.e.* "astro-particle in space") are a community eager of resources and poor in terms of man-power for computing: we're willing to test any solution to increase our pool of resources and to keep up with the software infrastructure developments, given the limited man-power available
- Given the nature of the partners for the various projects (ASI, Chinese collaborators, …) we can have small and/or temporary resources at our disposal: merging them in a single batch system would be a big added value
- As a new experiment in a design phase, HERD represents the perfect opportunity to migrate towards a deeper integration with INFN-Cloud provided services and the first impression is overwhelmingly positive.