A cloud-based computing infrastructure for the HERD cosmic-ray experiment

Diego Ciangottini, Matteo Duranti,
Valerio Formato, Nicola Mori, Daniele Spiga

Istituto Nazionale di Fisica Nucleare

ACAT 2022 – Bari (IT) - 25th October 2022

Introduction

- **HERD: space-based cosmic-ray direct detection experiment**
  - To be installed in 2027 on the forthcoming Chinese Space Station
- **International collaboration led by Chinese institutions with a strong European component (mostly Italian)**
- **Currently in advanced design and prototyping stage**
  - Possibility to start developing code and computing model from scratch
- **Facts about computing models for cosmic-ray experiments:**
  - Exp. “size” << LHC
    - “Easy”, but…
  - High demand of raw computing power in some scenarios
    - E.g. upper energy limit for MC simulations in the PeV region
  - Difficulties in exploiting some computing optimization techniques
    - E.g. optimized workloads like fast MC simulations may not be suitable for assessing the e/p rejection power with sufficient accuracy
    - Small community with relatively little manpower to devote to this topic
  - Highly dynamic use of computing resources (short-period bursts)
    - Due to e.g. small number of involved people leading to significant usage fluctuations

Introduction

- **The cloud can provide some solution:**
  - Add/remove resources with minimal overhead to dynamically cope with high/low demand periods
  - Efficiently exploit opportunistic resources with cloud-native solutions
  - Deploy self-hosted, self-managed services
- **The computing model must be designed to fully profit of the cloud**
  - E.g. to maximize the usage of opportunistic resources by minimizing the set-up period after the resources are made available in the cloud
- **Towards a fully-cloud based model for HERD:**
  - R&D and prototyping work
  - Based on INFN Cloud infrastructure and resources

Overview

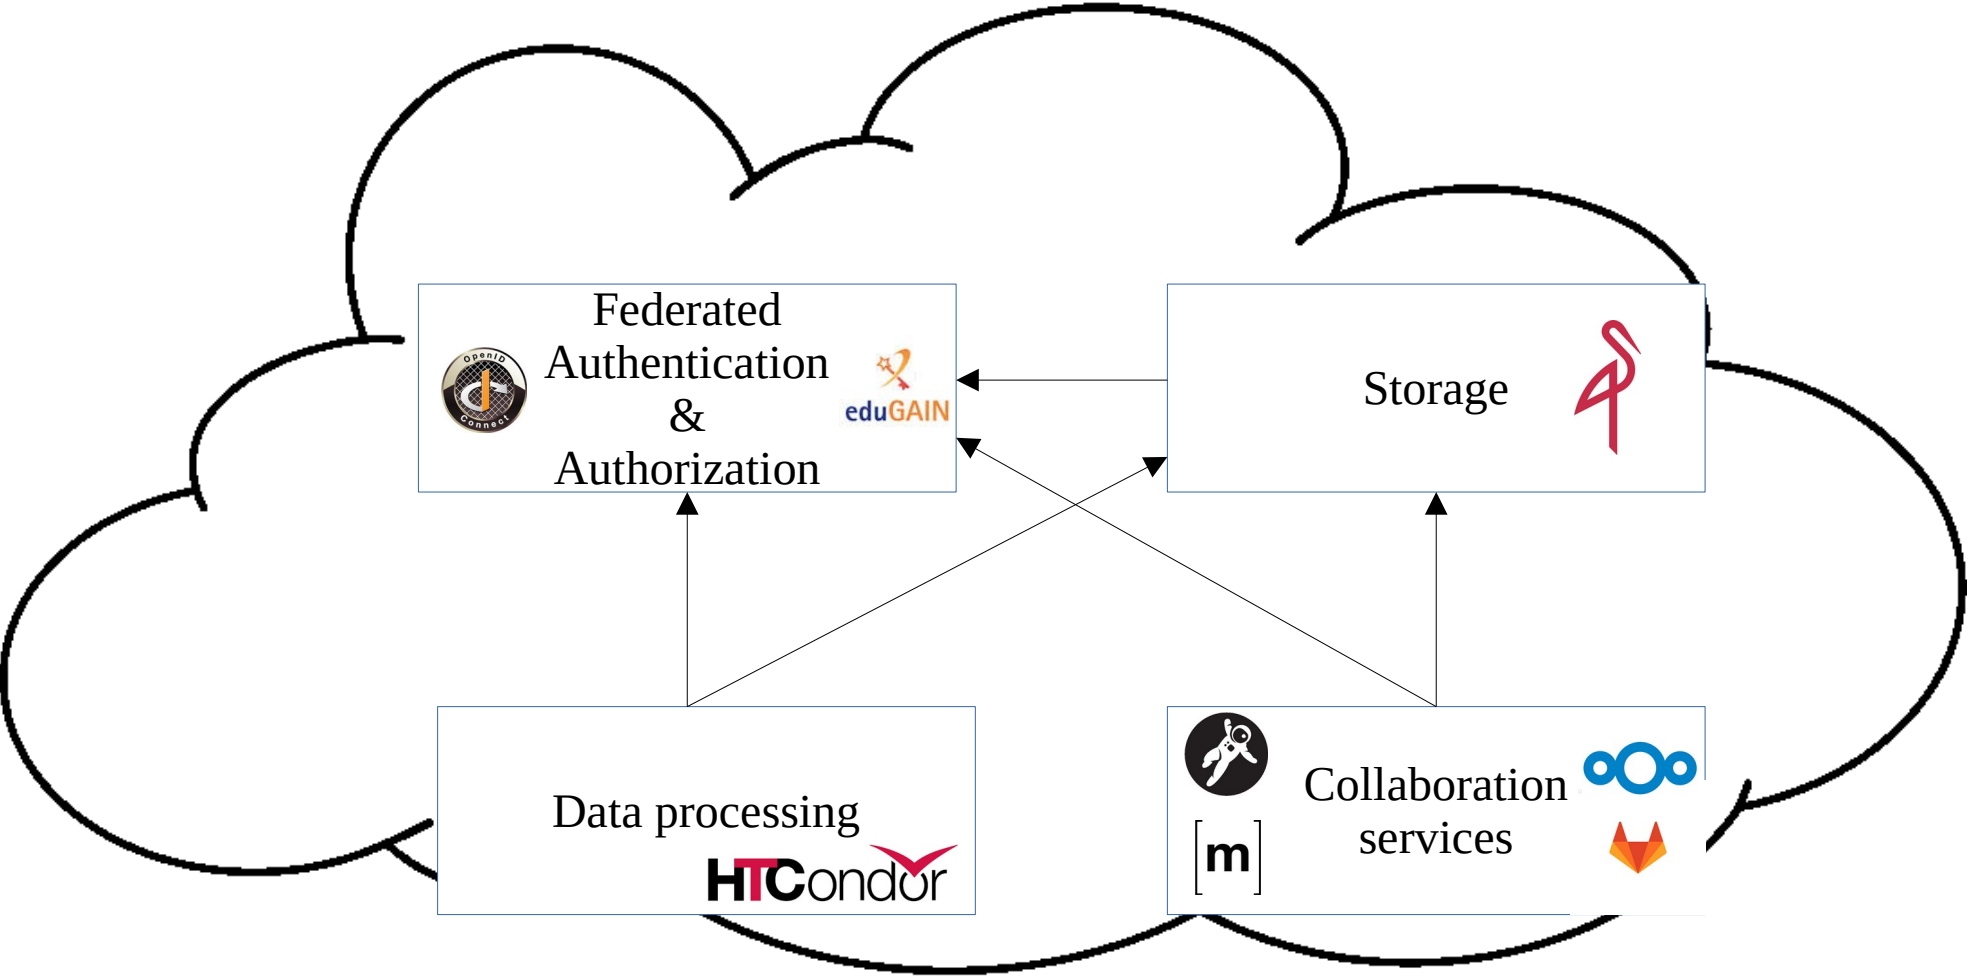
Federated
Authentication
&
Authorization
eduGAIN
Storage
Data processing
HTCondor
Collaboration
services

Overview

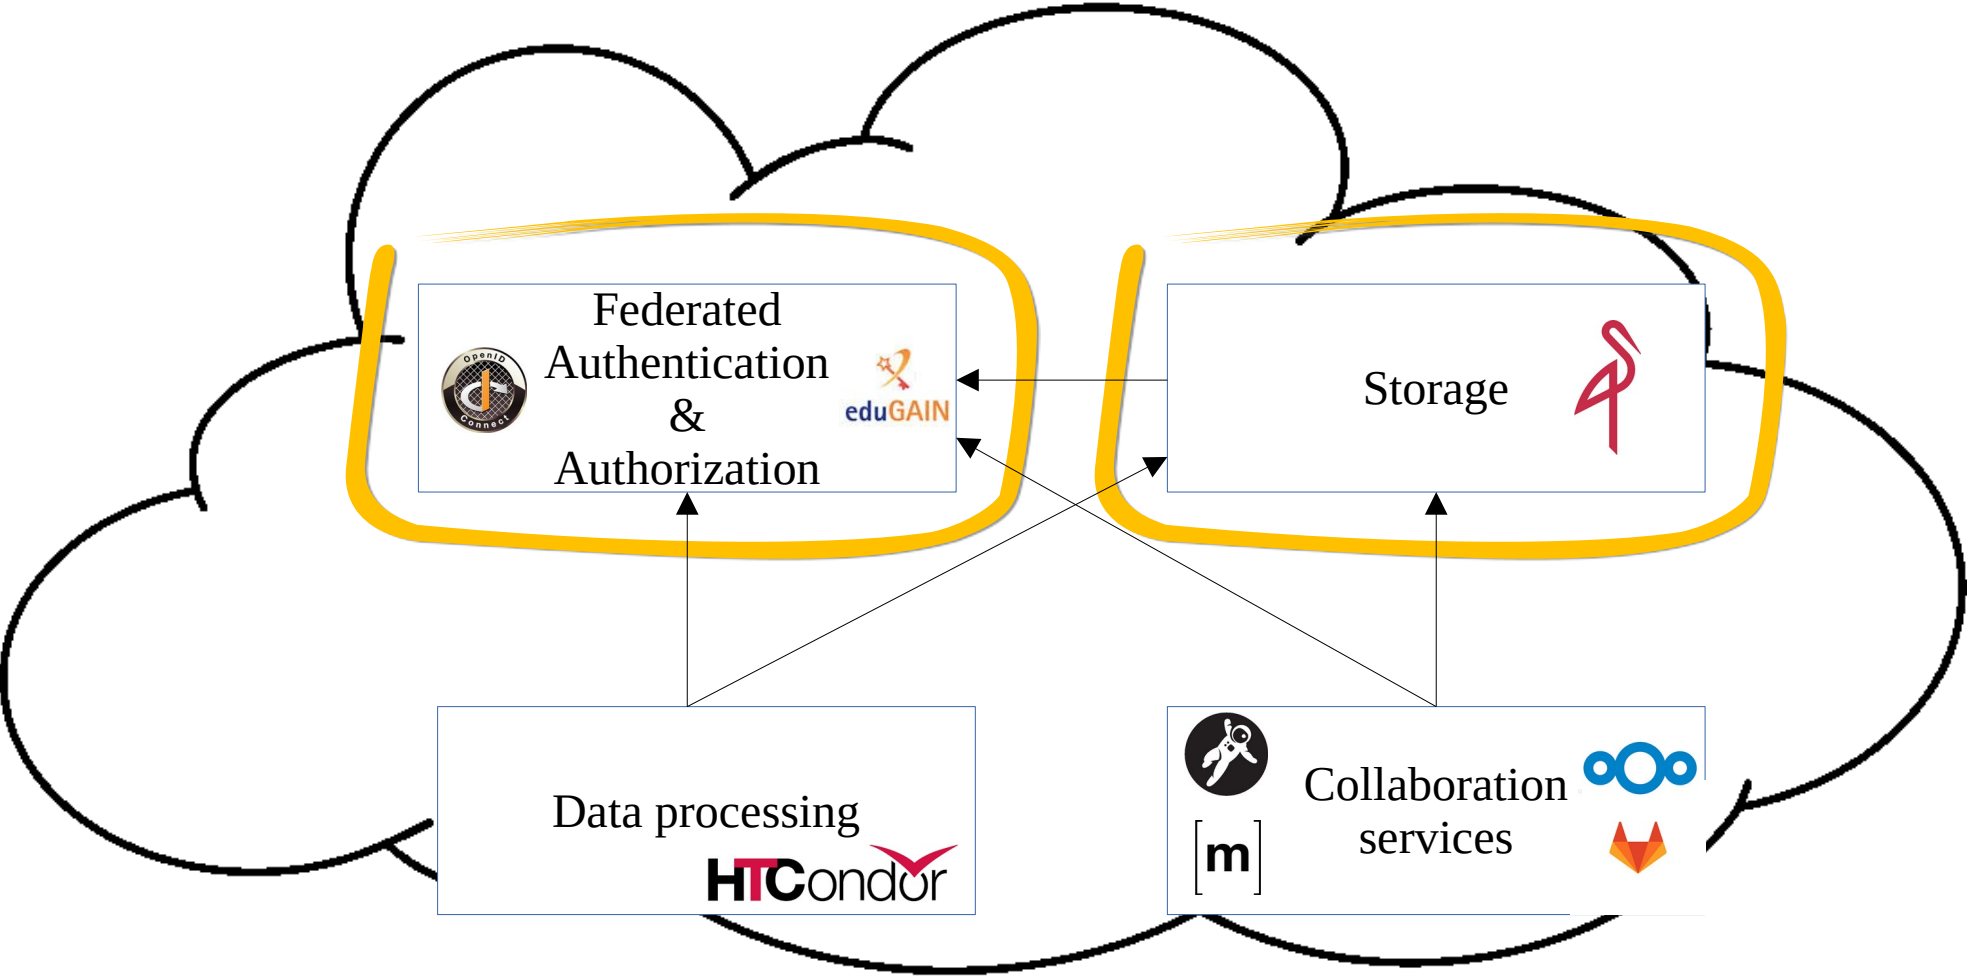
Federated
Authentication
&
Authorization
eduGAIN
Storage
Data processing
HTCondor
Collaboration
services

AuthN/Z

- **Single Sign On based on a dedicated INDIGO-IAM instance (IAM HERD)**
- **Federated with institutional SSO services directly and through EduGAIN**
- **OpenID Connect protocol**

AuthN/Z

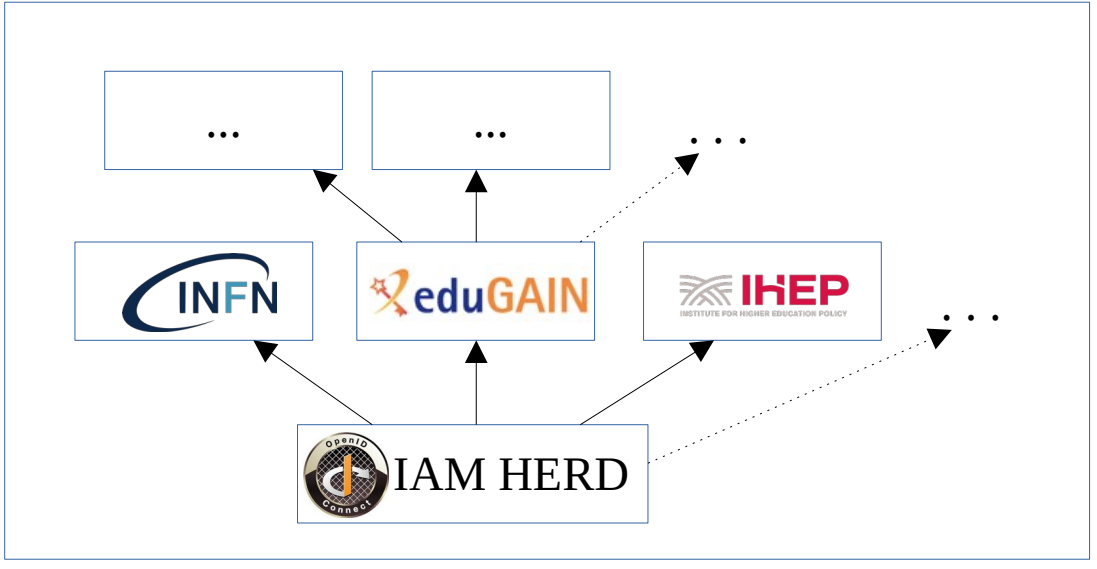

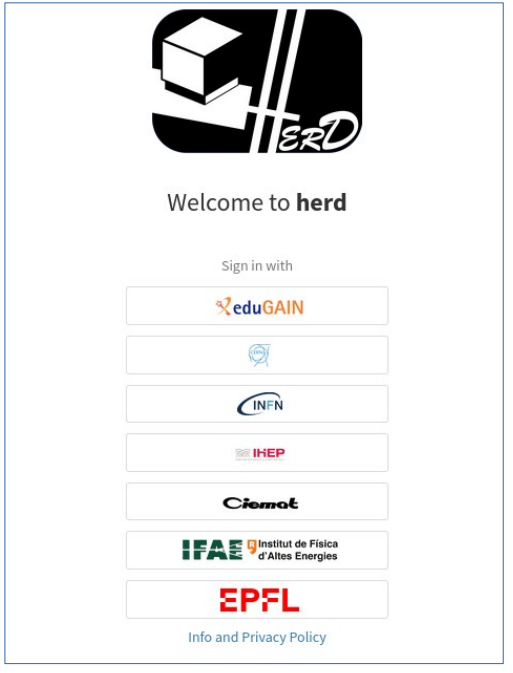

https://iam-herd.cloud.cnaf.infn.it

Storage

- **S3-compatible storage based on MinIO**
  - Well-established cloud storage technology
  - Available storage backends for many components (e.g. CVMFS)
- **Used for data, software and services**
- **Testbed based on 4 MinIO instances running on Kubernetes**
  - 3+1 erasure coded storage
  - 100 TB raw capacity provided via Ceph (DICE project, grant agreement ID: 101017207)
- **Access control:**
  - OIDC tokens managed via oidc-agent (users)
  - Access keys (services)

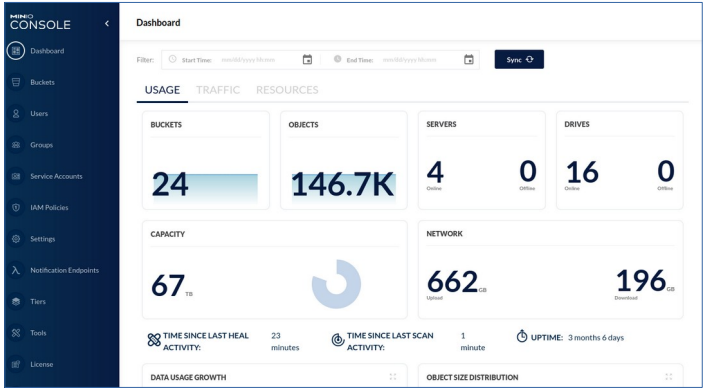

MinIO HERD dashboard

Data processing

- **On-demand HTCondor cluster running on Kubernetes**
- **Local containerized clients for submitting jobs**
  - Authentication on the cluster via OIDC tokens (oidc-agent)

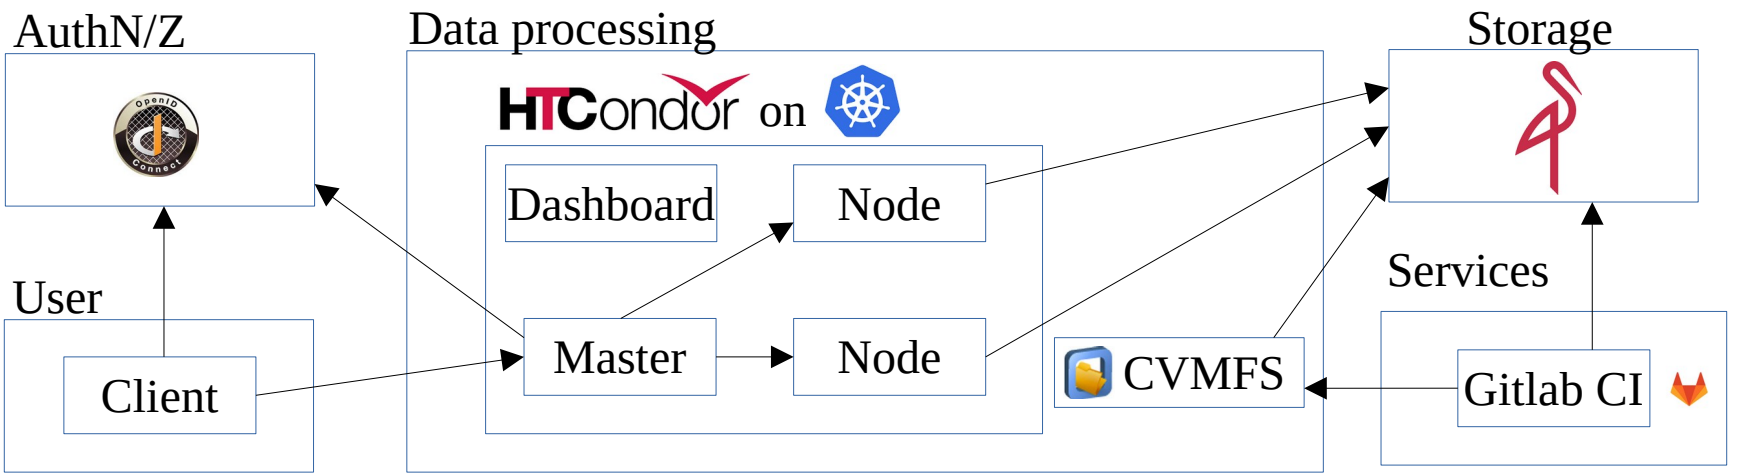

Data processing

- **Computing nodes S3 I/O via pre-signed URLs**
  - Retrieved at job launch using user's token and stored on the node local storage
- **SW distribution via CVMFS**
  - S3 backend
  - Automated deployment via Gitlab CI

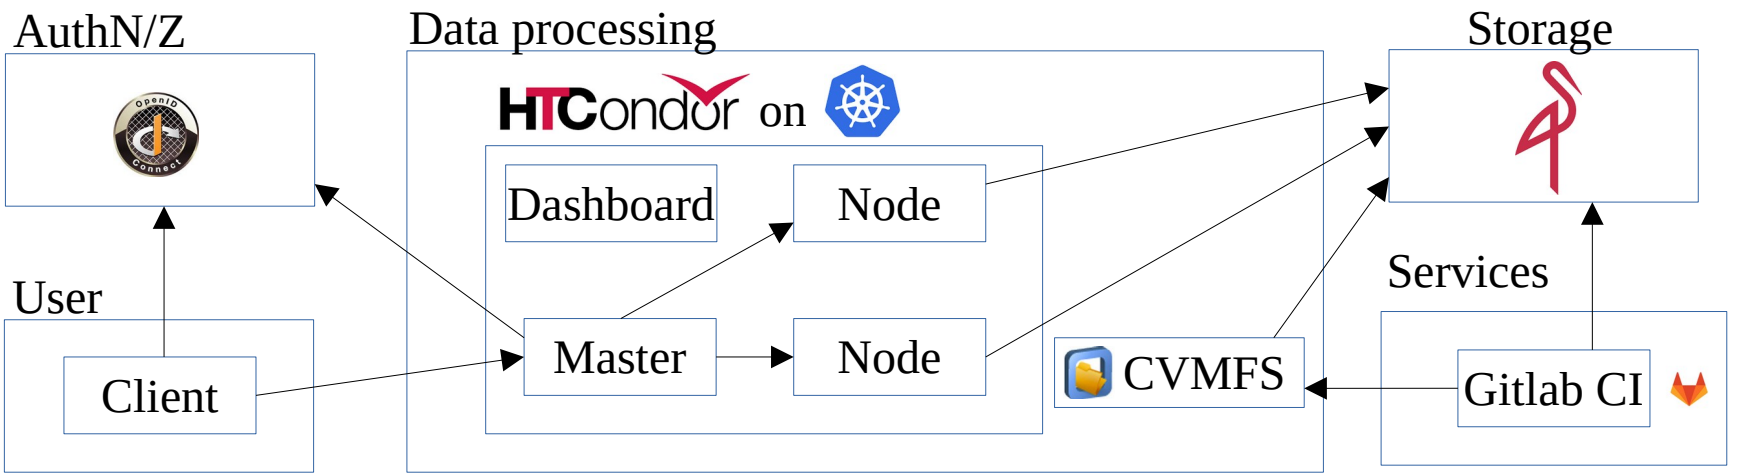

Data processing

- **Automated deployment of collaboration code to CVMFS via Gitlab CI**
  - Triggered by commit on master / new release tag (x.y.z)
  - Concurrent deployment for different OSs

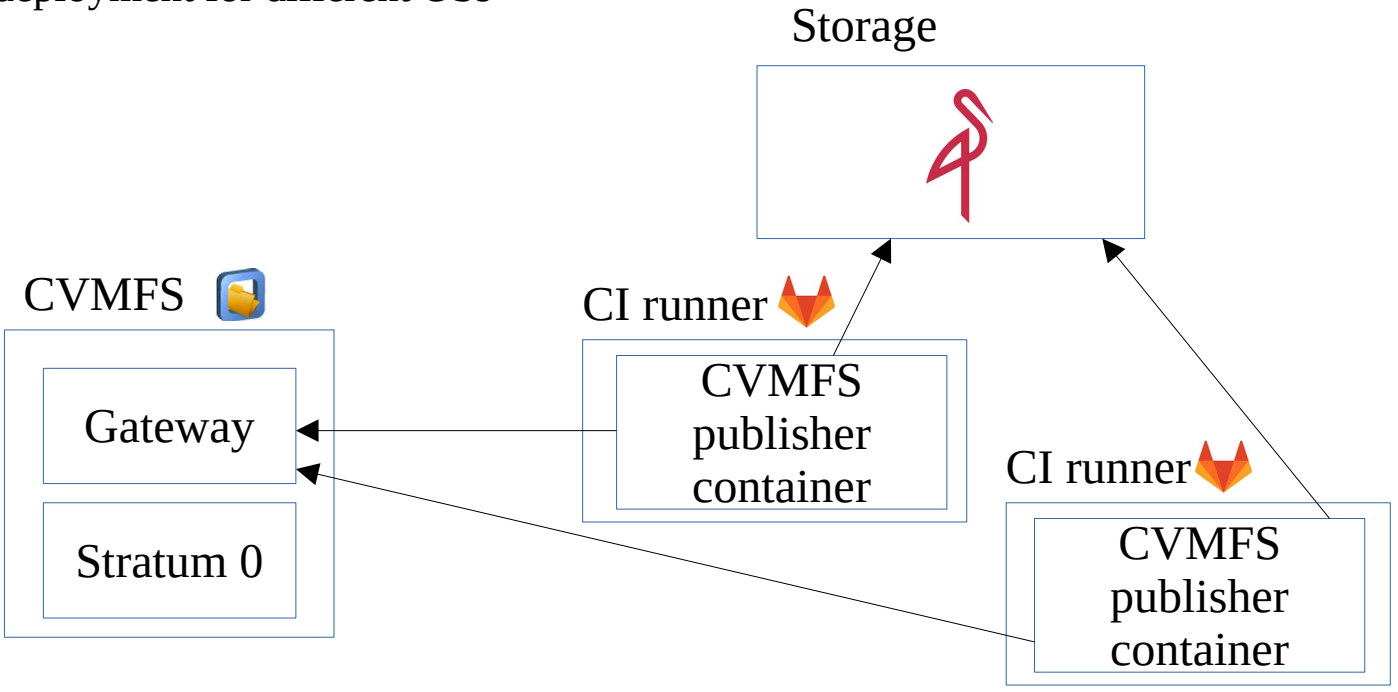

Data processing

- **The full HERD data processing chain has been successfully tested on the cloud infrastructure**
  - Launch script: retrieves the psURLs for all input and output files, prepares the Condor submit file and the job script, and submits the job cluster
  - Simulation: based on Geant4, output file is copied on S3 with curl using a psURL
  - Digitization: the job script reads the simulation input file from S3 using curl and a psURL, executes the digitization, and stores the digitized output file on S3 using a psURL
  - Reconstruction: same workflow as digitization, with a digitized file as input
  - Analysis: same workflow as reconstruction, with a reconstructed file as input
- **No major issues encountered**
- **Performance still to be investigated**
  - Sub-optimal storage testbed (MinIO on VMs with Ceph backend storage)
  - Single-site setup, no multi-site optimizations yet
- **Open issues**
  - S3 access control for long-queueing/long-running jobs
    - psURLs validity currently limited to 1 week
  - Distribution of user's code (CVMFS?)
    - Using HTCondor file transfer at the moment
  - Integration with "conventional" pledged resources (HTCondor flocking?)

Collaboration services

- **Peculiarities of the HERD collaboration:**
  - Cosmic-ray experiment → no umbrella organization like CERN → no institutional services available to every collaboration member
  - Chinese institutions → no access to free tools from e.g. Google
- **Solution: self-host the needed services at collaboration level**
  - The cloud approach ease this task
- **Fully-containerized web applications running on cloud VMs**
- **SSO through IAM HERD**
- **Automated backup/restore to/from S3**

Collaboration services

- **Self-hosted services:**
  - Web site with restricted area (Grav)
  - Calendar (NextCloud)
  - Document server (NextCloud)
  - Code repository (Gitlab)
  - CI/CD (Gitlab-CI runners)
  - Chat (Matrix , experimental)
  - Indico (maybe in future)
  - . . .
- **Open issues**
  - Secrets management
    - Currently hand-managed, planning to evaluate HashiCorp Vault for automated management
  - High availability (currently not an issue)

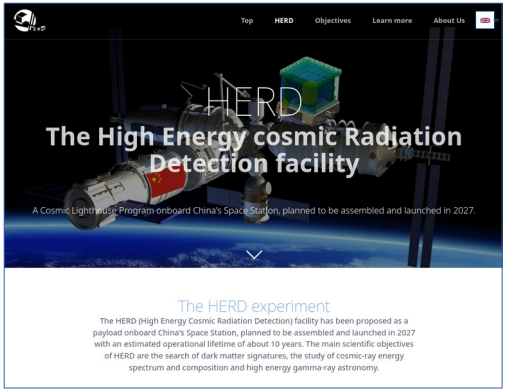

Web site

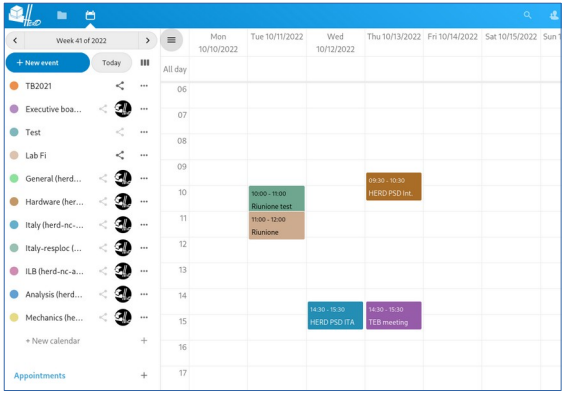

Calendar

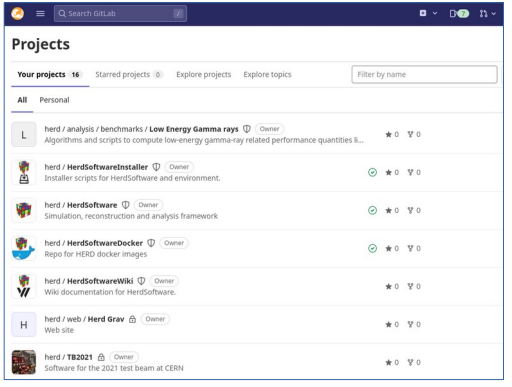

Code repo

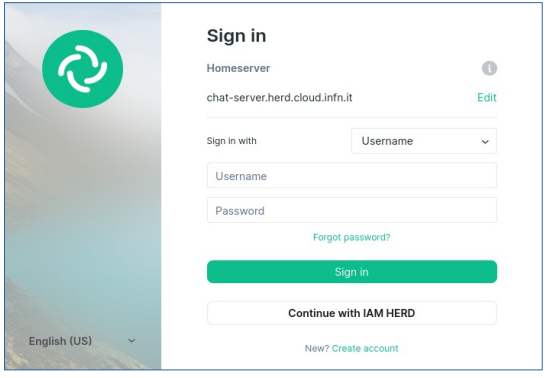

Chat

Summary

- **A fully cloud-based computation environment for the HERD cosmic-ray experiment has been successfully prototyped and tested**
- **The cloud can support a computing model that is more suited to cosmic-ray experiments than the pledge-based LHC ones**
- **The cloud offers the flexibility to self-host and manage services for the collaborations**
  - Even for low-manpower ones
- **Future developments**
  - AuthN/Z for long jobs
  - Distribution of user's code
  - Automated management of secrets
  - Federation with pledged resources
  - Usage of opportunistic resources

Thanks to the INFN Cloud staff for help and support!