# INFN-CLOUD-CNAF in EGI-ACE

Period 1 Review meeting, 24/05/2022

**Doina Cristina Duma**
Distributed Systems Manager – SDDS
INFN -CNAF

**Dissemination level:** Public/Confidential

**Disclosing Party:** Project Consortium

**Recipient Party:** European Commission

EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.

# Outline

- Who we are and what we do  - INFN & INFN CNAF
- INFN Cloud CNAF & EGI-ACE
- Integration points: Check-in, Accounting, AppDB, Monitoring
- Bringing GPUs to EGI
- Supported communities:
    - VIRGO, DigiFarm.io, FERMI-LAT, INACTIVE-SarsCov2, I-NERGY

# INFN (National Institute for Nuclear Physics)

- A long tradition in **state-of-the-art distributed IT technologies**, from the first small clusters to Grid and Cloud-based computing.
- INFN is not interested in computing per-se, but as an essential way to **support its research and mission**.
- For the past 10 years, this mainly meant supporting the experiments @ CERN (LHC), although the scope is now widening very quickly to other communities.
- Currently, INFN operates:
  - 9 medium size centers (Tier-2s in the LHC Computing Grid lingo)
  - 1 large Tier-1 center, at CNAF (Bologna)
- All the INFN centers are connected with 10-100 Gbit/s dedicated connections through the GARR network.
- Collectively, our main centers have about 65,000 CPU cores, 50PB of enterprise-level disk space, 60PB of tape storage.

# CNAF history & mission

1960 | 1970 | 1980 | 1990 | 2000 | 2010 | 2020

- **CNAF** (**C**entro **N**azionale **A**nalisi **F**otogrammi – National Center for Frames Analysis) is established in 1962 as an INFN Central Facility for the analysis of frames coming from bubble chambers.
- CNAF plans, develops and manages the INFN wide area network, which gradually evolved into the Italian research network, now managed by GARR (1980-2000)
- At the end of the 90s, CNAF realizes the LHC Italian Tier1 Data Center.
- CNAF becomes one of the main actors in the development of GRID World Wide Computing.
- CNAF has a solid Data Center, ISO-27001 certified.
- It offers CPU and storage resources to more than 30 INFN physics experiments, as well as to other disciplines.
- It develops innovative Cloud services oriented to the scientific world, to industry and society.
- It is a key player in technology transfer (through its TTLab) and is active in many projects of national or international relevance.

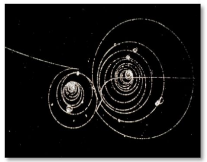

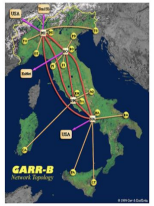





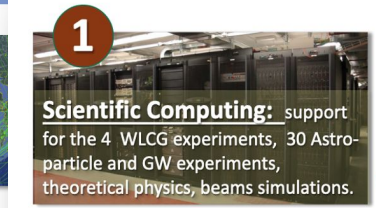

# INFN-CLOUD-CNAF

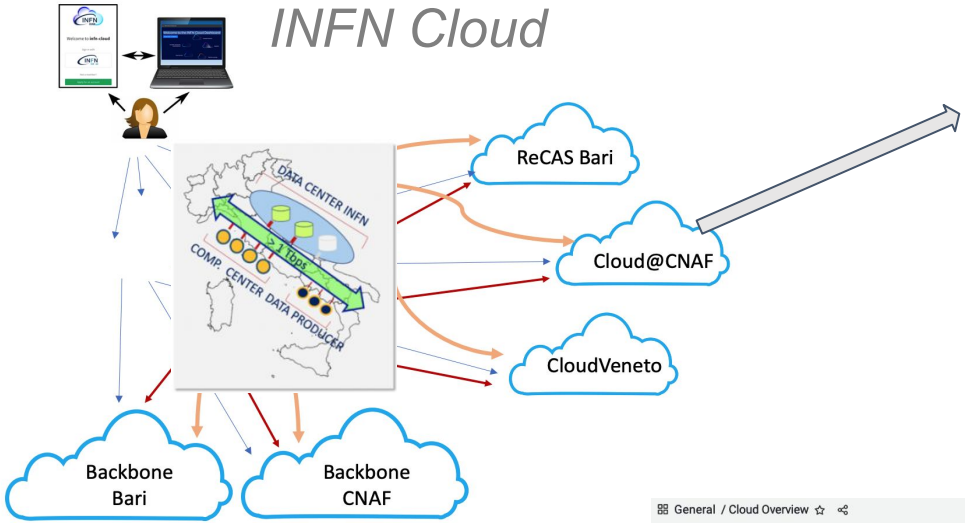


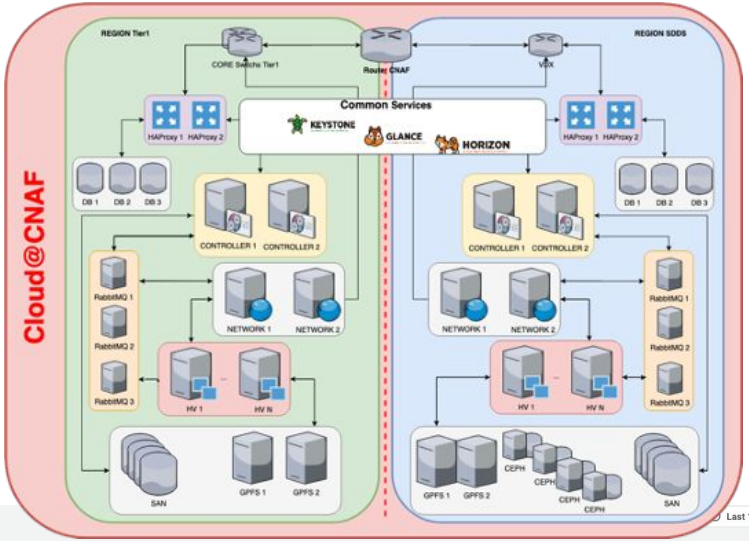



| | Quantity | UM | Total units | Units/month | Total M1-M15 |
|---|---|---|---|---|---|
| | | | | | |
| CPU | 200 | CPU core/hour | 4,380,000.00 | 146,000.00 | 520,426.00 |
| GPU | 2 | GPU server/hour | 43,800.00 | 1,460.00 | 8,765.00 |
| Storage | 165 | TB/month | 4,950.00 | 165.00 | 217.47 |

# INFN-CLOUD-CNAF in EGI-ACE, EGI Check-in

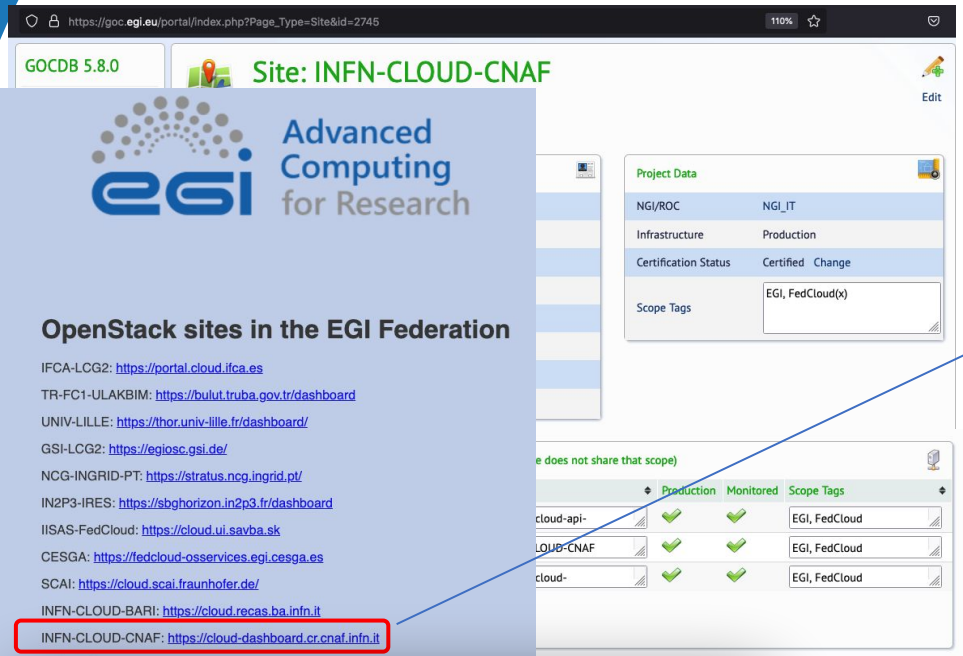


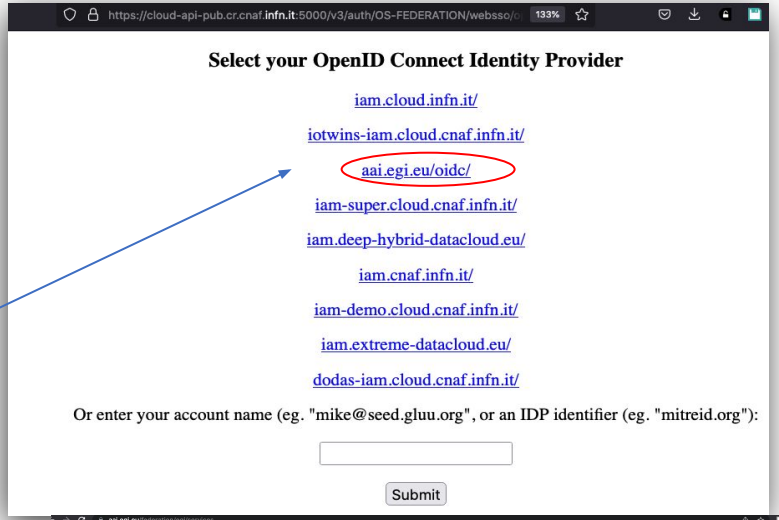


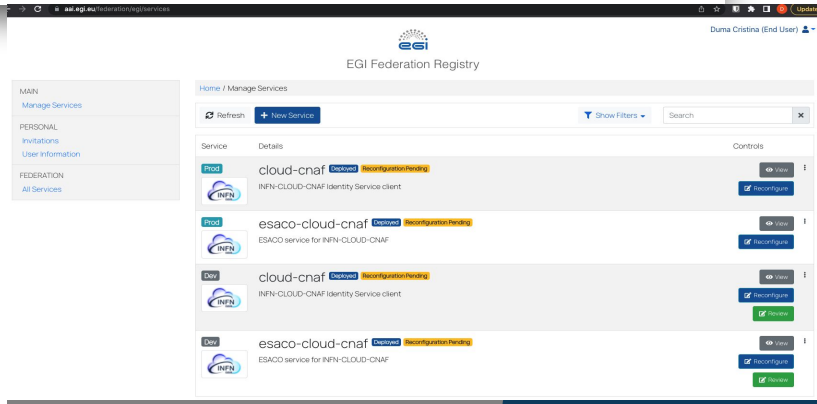


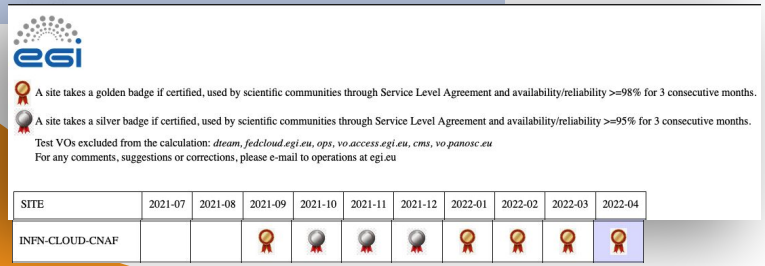

# Easy Integration - Accounting & Monitoring

https://accounting.egi.eu/

The Cloud Compute EGI view shows the accounting data in Resource Centre INFN-CLOUD-CNAF. The metric shown is Elapsed time * Number of Processors, grouped by VO and Month, all VOs are shown.

Resource Centre INFN-CLOUD-CNAF — Elapsed time * Number of Processors (hours) by VO and Month (All VOs)

| VO | Jan 2022 | Feb 2022 | Mar 2022 | Apr 2022 | May 2022 | Total | Percent |
|---|---|---|---|---|---|---|---|
| digifarm.io | 47,616 | 43,008 | 47,616 | 6,720 | 0 | 161,543 | 19.79% |
| dteam | 1,488 | 2,008 | 0 | 0 | 0 | 14,914 | 1.83% |
| fermi-lat.infn.it | 0 | 0 | 23 | 0 | 785 | 808 | 0.1% |
| ops | 20 | 14 | 25 | 22 | 8 | 160 | 0.02% |
| virgo | 38,688 | 34,944 | 112,066 | 115,200 | 63,360 | 552,087 | 67.62% |
| vo.i-nergy.eu | 0 | 0 | 0 | 31,306 | 31,680 | 62,986 | 7.71% |
| vo.inactive-sarscov2.eu | 0 | 0 | 0 | 11,255 | 12,672 | 23,927 | 2.93% |
| Total | 87,812 | 79,973 | 159,731 | 164,503 | 108,505 | 816,424 | |
| Percent | 10.76% | 9.80% | 19.56% | 20.15% | 13.29% | | |

1 - 7 of 7 results

Download JSON Data / Download CSV Data

https://argo.egi.eu

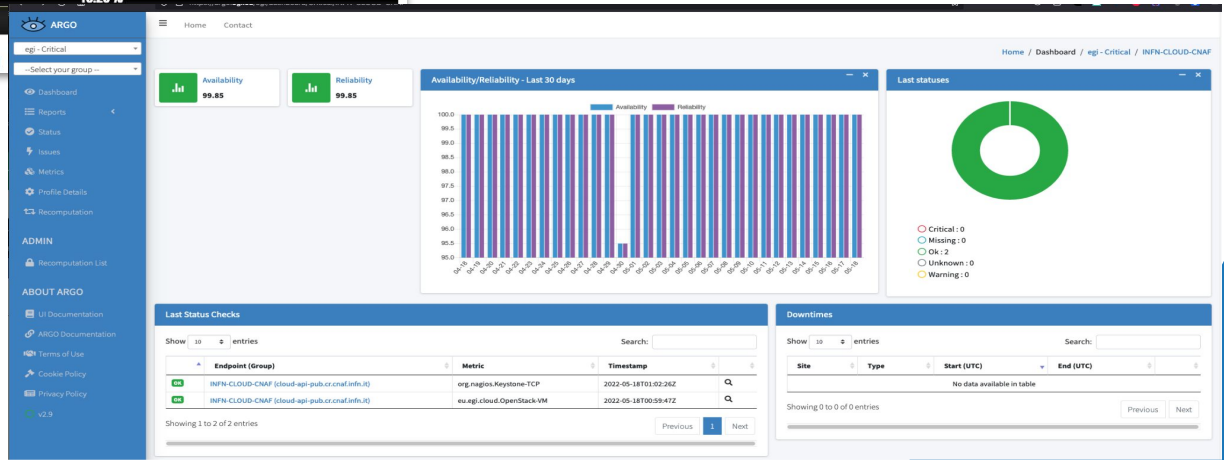

# Bringing GPUs

- Supporting use-cases that needed to exploit specialized hardware like GPUs:
  - DigiFarm.io & I-NERGY : NVIDIA V100
  - Amber INACTIVE SarsCov2: NVIDIA A100

- Defining different types of “flavors” to accommodate the specific requirements:

| Name | RAM | Disk | VCPUs |
|---|---|---|---|
| df.16CPU_32GB_100GB_1V100 | 32768 | 100 | 16 |
| df.16CPU_64GB_100GB_1V100 | 65536 | 100 | 16 |
| df.8CPU_16GB_100GB_1V100 | 16384 | 100 | 8 |
| df.8CPU_32GB_100GB_1V100 | 32768 | 100 | 8 |
| | | | |
| ina.16CPU_64GB_160GB_1A100 | 65536 | 100 | 16 |
| ina.16CPU_64GB_160GB_2A100 | 65536 | 160 | 16 |
| | | | |
| inergy.16CPU_64GB_100GB_1V100 | 65536 | 100 | 16 |

# VIRGO

The relevant services to be deployed for the VIRGO (and more generally IGWN) low-latency alert generation infrastructure are:

- The Gravitational-Wave Candidate Event Database (**GraceDB**): it provides a centralized location for aggregating and retrieving information about candidate gravitational-wave events.
- The LIGO-Virgo Alert Network (**LVAlert**): a prototype notification service
- **GWCelery**: a service for annotating and orchestrating IGWN alerts

**Successfully managed to run a test tier of the GraceDB and LValert servers, according to their plans**

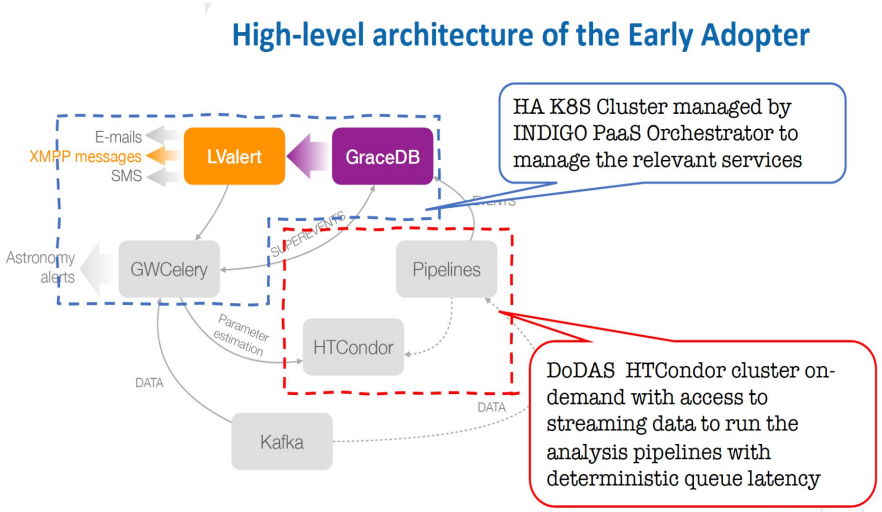


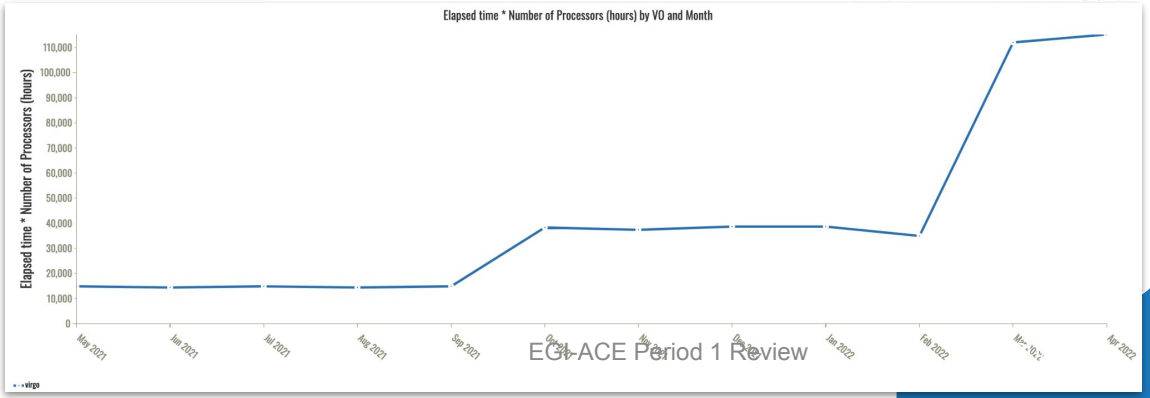

# FERMI-LAT & DODAS

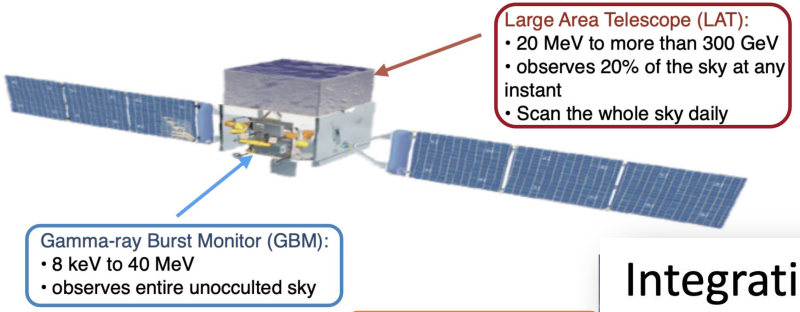


The **Fermi Gamma-ray Space Telescope** - space observatory being used to perform gamma-ray astronomy observations from low Earth orbit.

- **Large Area Telescope (LAT),** main instrument used to perform an all-sky survey studying **astrophysical and cosmological phenomena** such as active galactic nuclei, pulsars, other high-energy sources and dark matter

## Integrating DODAS and Fermi-LAT

DODAS has been used to provide a all-in-one system with

- HTCondor batch on-demand
  - Support Token based Authentication to allow remote job submission
  - User tailored runtime environment
  - Ready to support a cluster federation
- MinIO as a Service
  - Deployed as MinIO Kubernetes Operator
  - Fully integrated with runtime environment ( read and write )
  - Ready to support replicas to distributed clusters

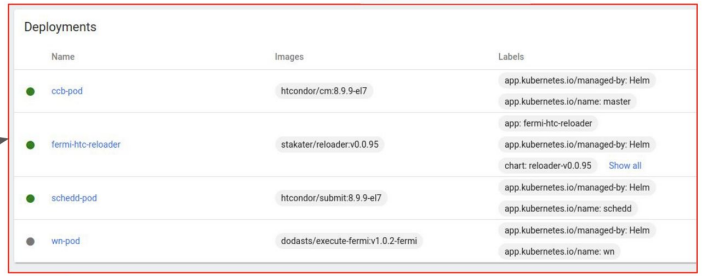


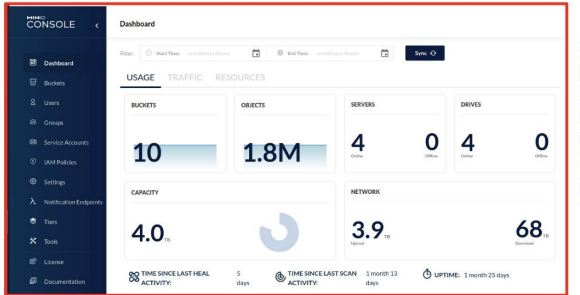

# AMBER-based modelling of SARS-CoV-2 Spike protein

**Scientific impact:** Results obtained through the implementation of the proposed use case will be the basis for the implementation of specific research studies aimed at **identifying antiviral drugs and anionic polymers targeted to inactivate the Spike protein**

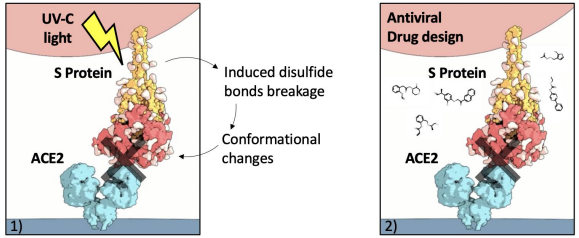


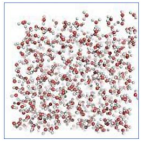

**Analysis**

- ✓ Molecular dynamics (MD) simulations of 500ns (x3) of classic MD and 500ns (x2) of Gaussian accelerate (MD) of spike protein wild-type and mutant are running.
- ✓ total atoms number of the system: 500000 (glycoprotein + ions + water)

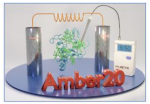


**Software**

- ✓ Amber is used for performing the molecular dynamic simulations
- ✓ It is the fastest academic GPU molecular dynamic simulations engine
- ✓ The software license is provided by CNR-ITB

**Computing Resources**

- ✓ Modern GPU devices based on the Ampere architecture able to speed-up the molecular dynamic simulations
  - ➢ N. 4 A100 GPUs made available by INFN
  - ➢ N. 4 A40 GPUs made available by CESNET

**Storage space**

- ✓ A large amount of storage space is required to store molecular dynamic simulations of raw outputs
  - ➢ 60 TB made available by INFN
  - ➢ 170 TB made available by CESNET

- Running Dynamics on A100
  - cMD(500ns) x 3
  - GaMD(500ns) x 2
  - 5 MDs x 5 Tb = 25 TB
- Running Dynamics on A40
  - 8 complexes x 3 replica =24 MDs
  - 24 MDs x 1 TB -> 24 TB

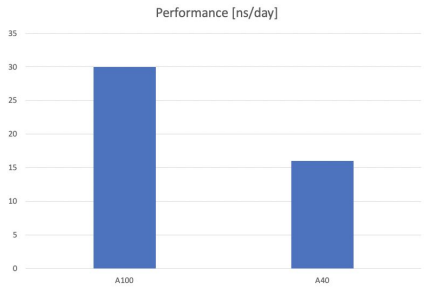


- **Dynamics are running on VMs instanced on Openstack equipped with**

| | |
|---|---|
| **VCPUs** | 16 |
| **RAM** | Min: 64GB – Max: 120GB |
| **Size** | Min: 80GB – Max: 160GB |
| **GPU** | 1 or 2 devices |
| **Volume** | Min: 20TB – Max: 80TB |

openstack.

- **Amber 20 has been dockerized by CNR-ITB to easily install and run it on the VMs.**
- **VMs have been properly configured to run Docker containers leveraging NVIDIA GPUs**

# I-NERGY - Artificial Intelligence for next generation energy services across Europe

Deliver an energy-specific **open modular framework for supporting AI-on-Demand in the energy sector (AI4 Energy)**

**Data Streaming:** Low latency near real time in-memory processing

**Data Storage:** According to the needs of services and tools

**Data Harmonisation:** Homogenisation and data pre-processing

**Data Ingestion:** Integration from heterogeneous sources

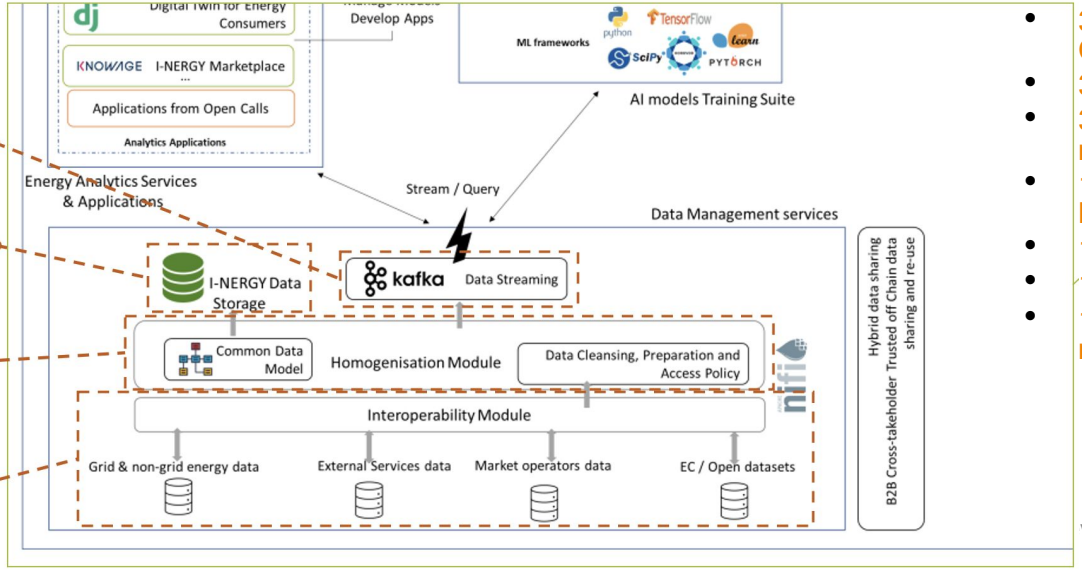


**Cloud Compute**:

- 3 VMs for APACHE NiFI Cluster (throughput target: 100 MB/s and 10000 events/s)
- 3 VMs for Confluent Kafka Cluster
- 3 VMs for MongoDB installation
- 3 VMs for Blockchain based notarisation
- 1 VM for evaluating ML/DL Models
- 1 VM for serving ML models
- 1 VM for the marketplace
- 1 VM with GPUs for training AI models (>30 GB of GPU RAM)

view

# DigiFarm.io

- Norwegian based ag-tech startup established in 2019.
- **Core vision** is to detect the world's most accurate field boundaries and seeded acres to power precision agriculture.
  - leveraging the latest advancements in Artificial Intelligence technology
  - developing deep neural network models for automatically detecting field boundaries through super-resolving Sentinel-2 satellite imagery to 1 meter resolution

## Resource Centre INFN-CLOUD-CNAF — Elapsed time * Number of Processors (hours) by VO and Month (Custom VOs)

| VO | Dec 2021 | Jan 2022 | Feb 2022 | Mar 2022 | Apr 2022 | Total |
|---|---|---|---|---|---|---|
| digifarm.io | 16,583 | 47,616 | 43,008 | 47,616 | 6,720 | 161,543 |

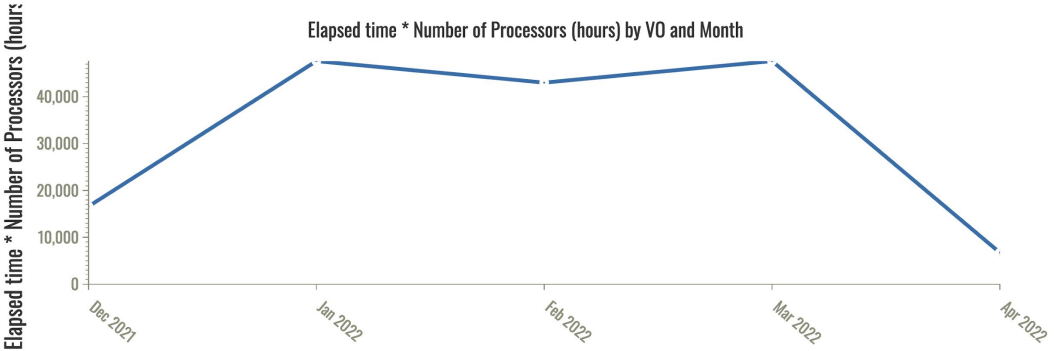

# Future

- Support chosen use cases until the end of their activities or end of EGI-ACE project
  - Colelct feedback on quality of resources offered
  - Provide feedback on EOSC marketplace solutions

- Fulfill commitments to the project
  - add support to new use-cases

# Thank you!

Contact: **egi-ace-po@mailman.egi.eu**
Website: **www.egi.eu/projects/egi-ace**

**EGI Foundation**

**@EGI_eInfra**

EGI-ACE receives funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 101017567.