www.egi.eu
@EGI_eInfra

EGI-ACE Overview

EGI Advanced Computing for EOSC

https://www.egi.eu/projects/egi-ace/

Gergely Sipos (EGI Foundation)
EGI-ACE Technical Coordinator

The work of the EGI Foundation
is partly funded by the European Commission
under H2020 Framework Programme

OpenAIRE-Nexus Launch event
10th March 2021
https://www.openaire.eu/openaire-nexus-public-launch-event

EGI-ACE objectives

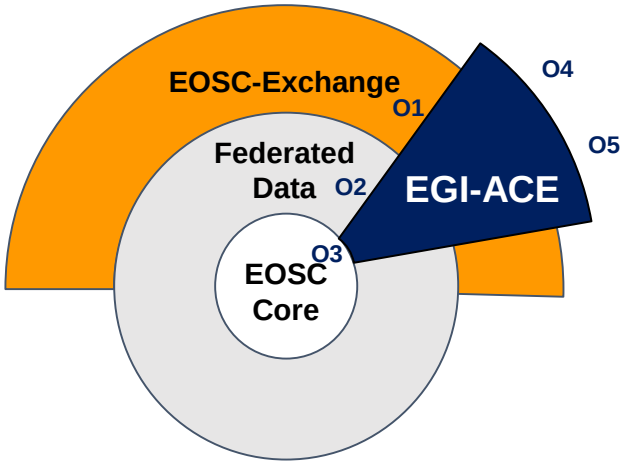

Objective 1. Deliver the European Open Science Cloud Compute Platform and expand the supply-side

Objective 2: Contribute to the implementation of the EU Data Strategy and the EOSC Data Commons to support the Green Deal, Health and Fundamental Research

Objective 3: Integrate the EOSC Compute Platform with the EOSC Portal and the EOSC Core

Objective 4: Contribute to the realization of a global Open Science Cloud

Objective 5: Expand the demand-side of EOSC across sectors and disciplines

EOSC Architecture:

Solutions for a sustainable EOSC

report from the EOSC Sustainability WG

EGI-ACE Consortium Overview

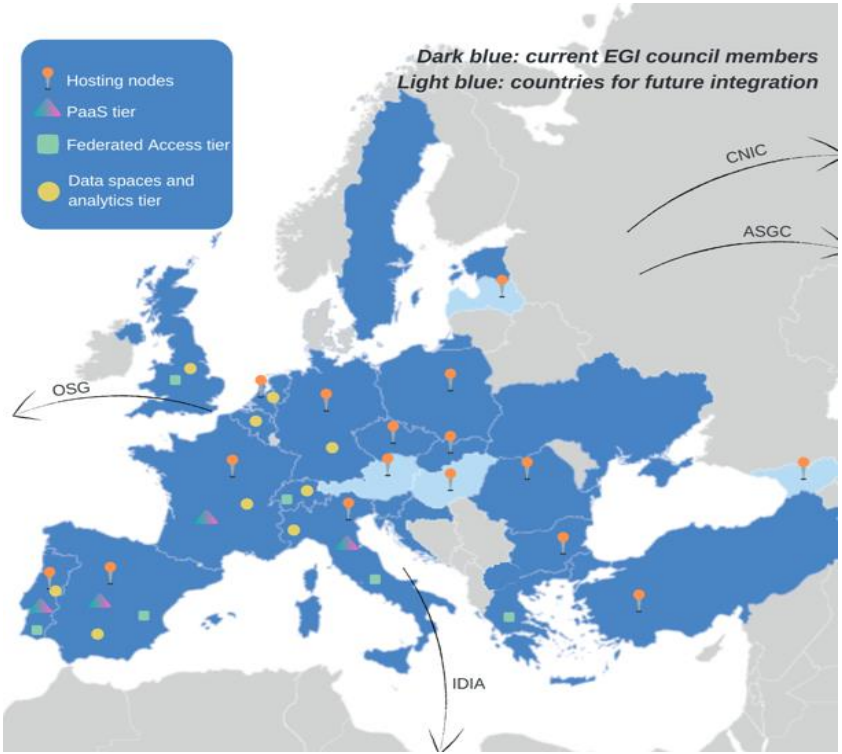

- EGI Foundation as coordinator
  - 33 participating partners, 23 third-parties
- Consortium at a glance
  - 16 partners to deliver cloud, HTC and HPC resources
  - 12 user access and platform solution providers
  - 21 providers from 13 communities to deliver data spaces
  - 12 federation service providers

  \+ Early adopters; Partner infrastructures

57% of the budget for service provisioning

EGI-ACE concept and methodology: Tier service architecture

Data Spaces and Analytics
Data and thematic data analytics and processing tools

Platforms
generic added-value platform level services

Federated Access
Federation-wide management of data and computing

Federated Resources
Compute and storage facilities

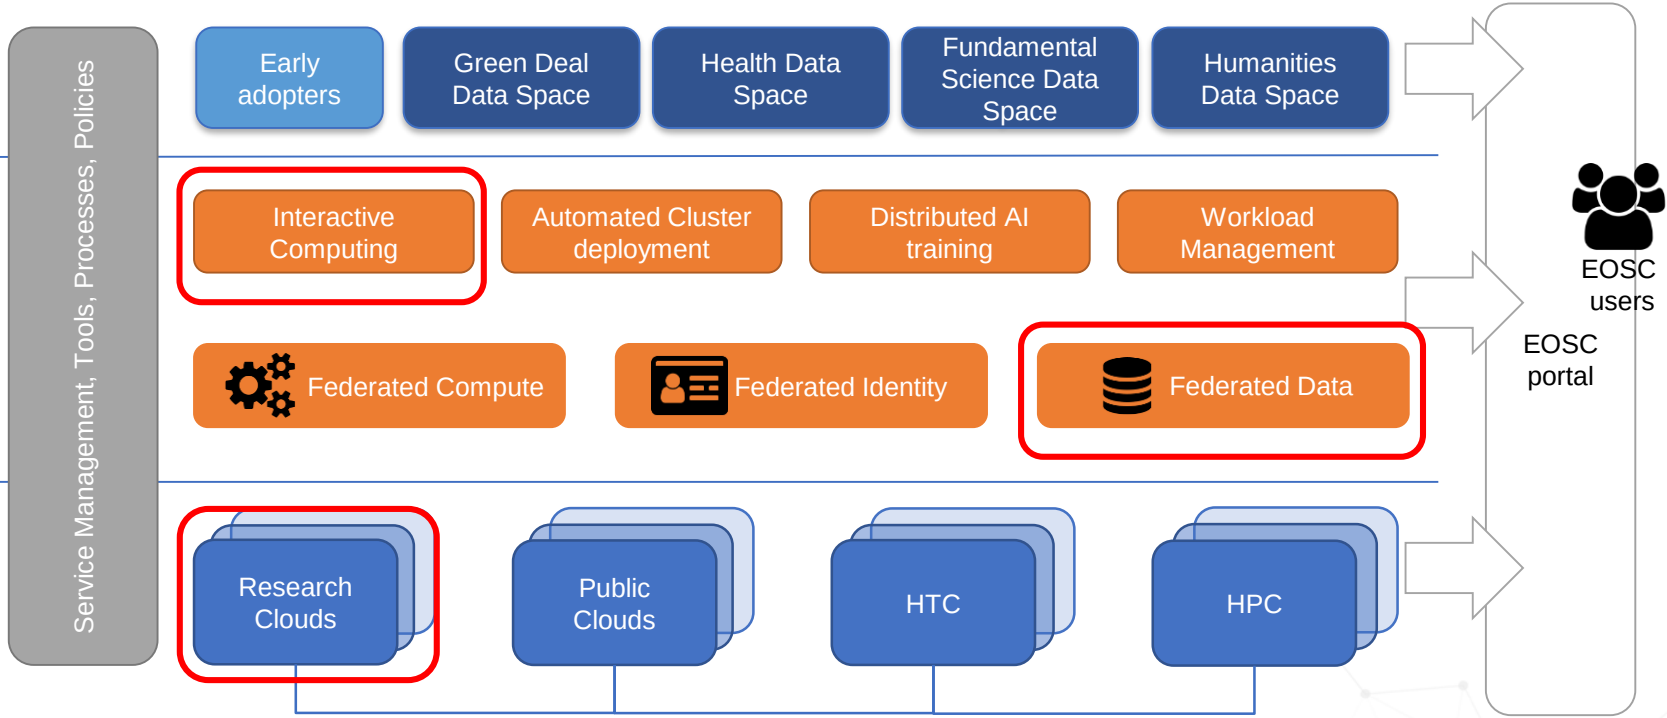

Example:
Reproducible big data analytics with notebooks

Notebooks + Binder + DataHub provide the means to reproducible Open Science with big data

Integrates with zenodo for publishing and discovery of notebooks

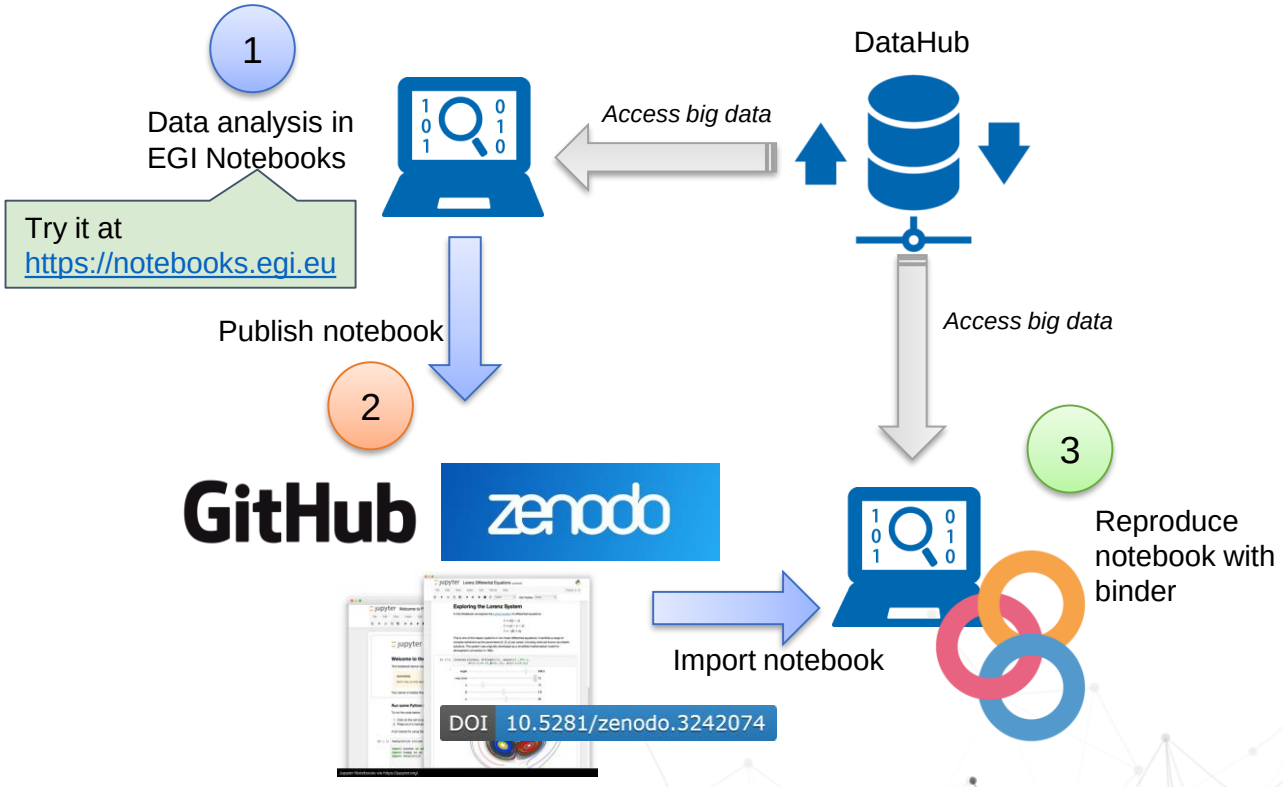

Compute services	Platform services	Data space services	Federation services (AAI, data, compute)
IM AppDB EGI HTC & Storage dCache Spider Storage Data Processing Compute EGI Cloud Compute AppDB DynDNS MetaCentrumCloud - CPU MetaCentrumCloud - GPU MetaCentrumCloud - Storage SCAI FedCloud v2 EGI - GSIOS IN2P3-IRES-CPU IN2P3-IRES-Storage TR-FC1-ULAKBIM - CPU TR-FC1-ULAKBIM - Storage INFN-BARI-CPU INFN-BARI-Storage INFN-CNAF-CPU INFN-CNAF-GPU INFN-CNAF-Storage INCD-Lisbon (NCG)-CPU INCD-Lisbon (NCG)-Storage EGI – IISAS DESY-FedCloud CESGA-CPU CESGA-Storage IFCA-LCG2-CPU IFCA-LCG2-Storage INCD-LIP-CPU INCD-LIP-Storage CYFRONET-CLOUD-CPU CYFRONET-CLOUD-Storage IICT-BAS-CPU IICT-BAS-Storage CLOUDIFIN-CPU CLOUDIFIN-Storage	EGI Notebooks & Binder EGI Workload Manager DIRAC DEEP training facility CSIC DEEP training facility LIP DEEP training facility DODAS	WeNMR portals UU WeNMR Portals UseGalaxy.eu Virtual Imaging Platform OpenRiskNet/NanoCommons Virtual Env. ENES Data Space OpenCoastS Prominence SURFSARA LOFAR Science Products NWO-I LOFAR Science Products Disaster mitigation and agriculture SeaDataNet WebOcean Data Analysis Iberian GBIF data space CSIC GBIF Cloud data space LIP GBIF Cloud data space EMSO ERIC data service OPERAS Metrics service OPERAS Certification service iop	EGI Check-in PERUN MasterPortal (EGI - Check-in) EC3 Indigo PaaS Orchestrator Rucio CVMFS OpenRDM EGI content distribution CVFMS EGI Data Transfer FTS EGI DataHub Onedata Rucio

31 different services
from
61 providers

EGI-ACE portfolio in the EOSC Marketplace

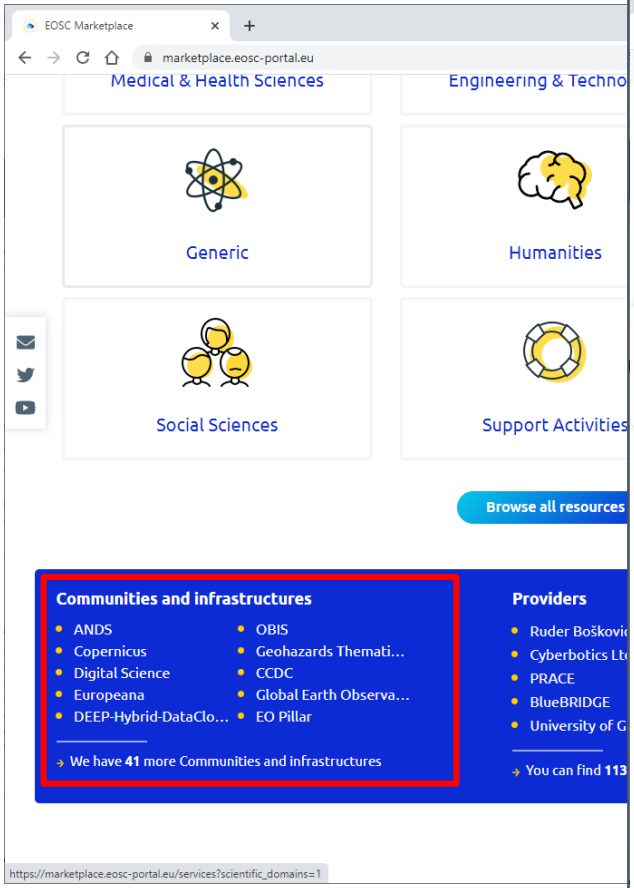

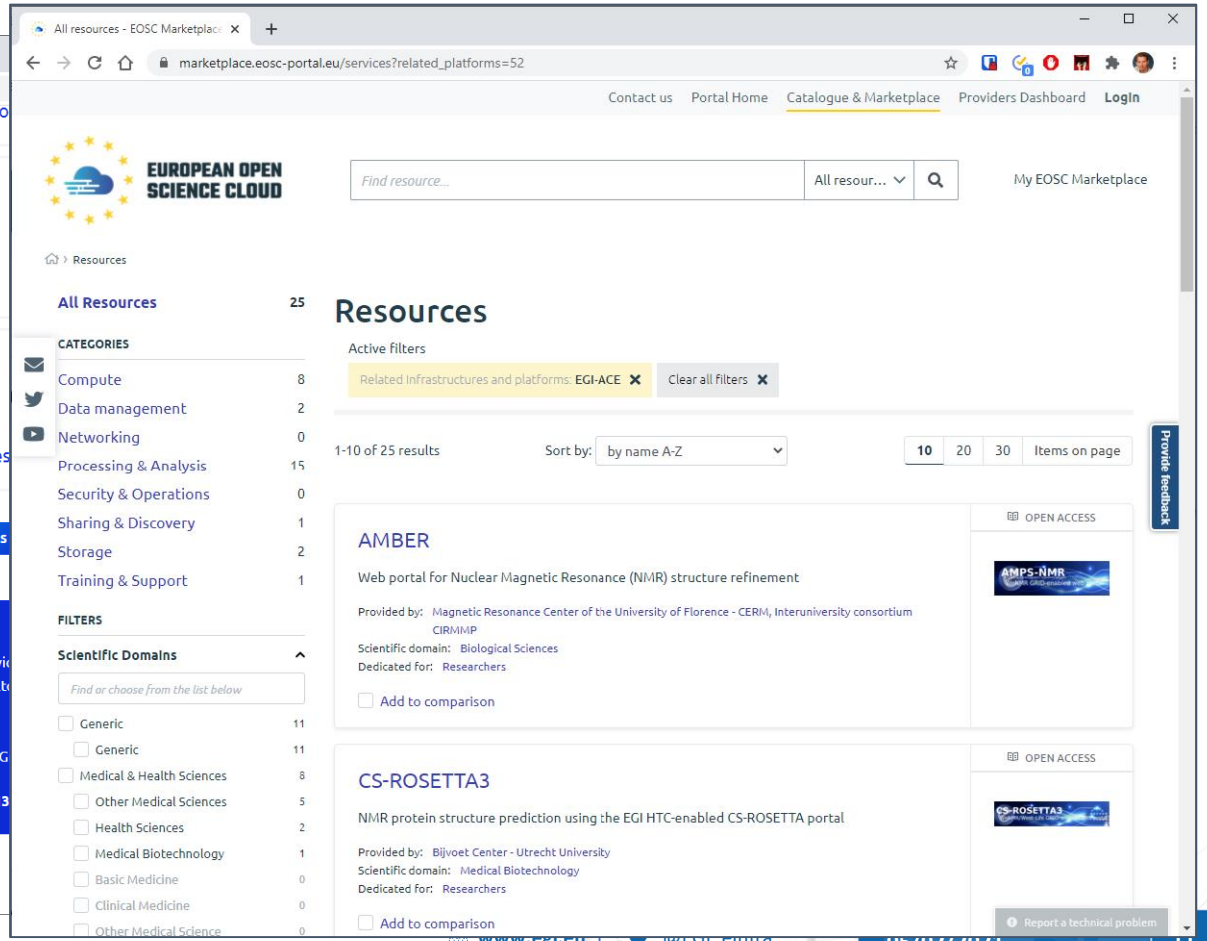

Infrastructure layer – Cloud, HTC, HPC

For users and communities within the project **AND** for new EOSC users

- **Cloud-HTC compute and short-mid term storage**:
  - 80 Million CPU hours / 250,000 GPU hours / 20 PB storage

- **Mix of funding mechanisms:**
  - Virtual Access:
  - Local grants enabling policy-based access (EGI Cloud Federation + new partners):
    - Moldova
    - Latvia
    - Georgia
    - Hungary
    - Austria
    - USA
    - China
    - South Africa
  - Pay-for-use:
  - Pilot HPC sites:

Diverse locations AND compute flavors

Resource allocation, user support (2/1)

Single users
Small groups
Experimental users

International projects
Multi-national communities

Business-to-User

EOSC Portal

- Ready-to-use resources/services
- Self-service configuration
- Short term engagement

Business-to-Business

EGI-ACE Open calls (To be published)
Evaluation 3x per year

- High capacity demand
- Custom configurations
- Long term engagement

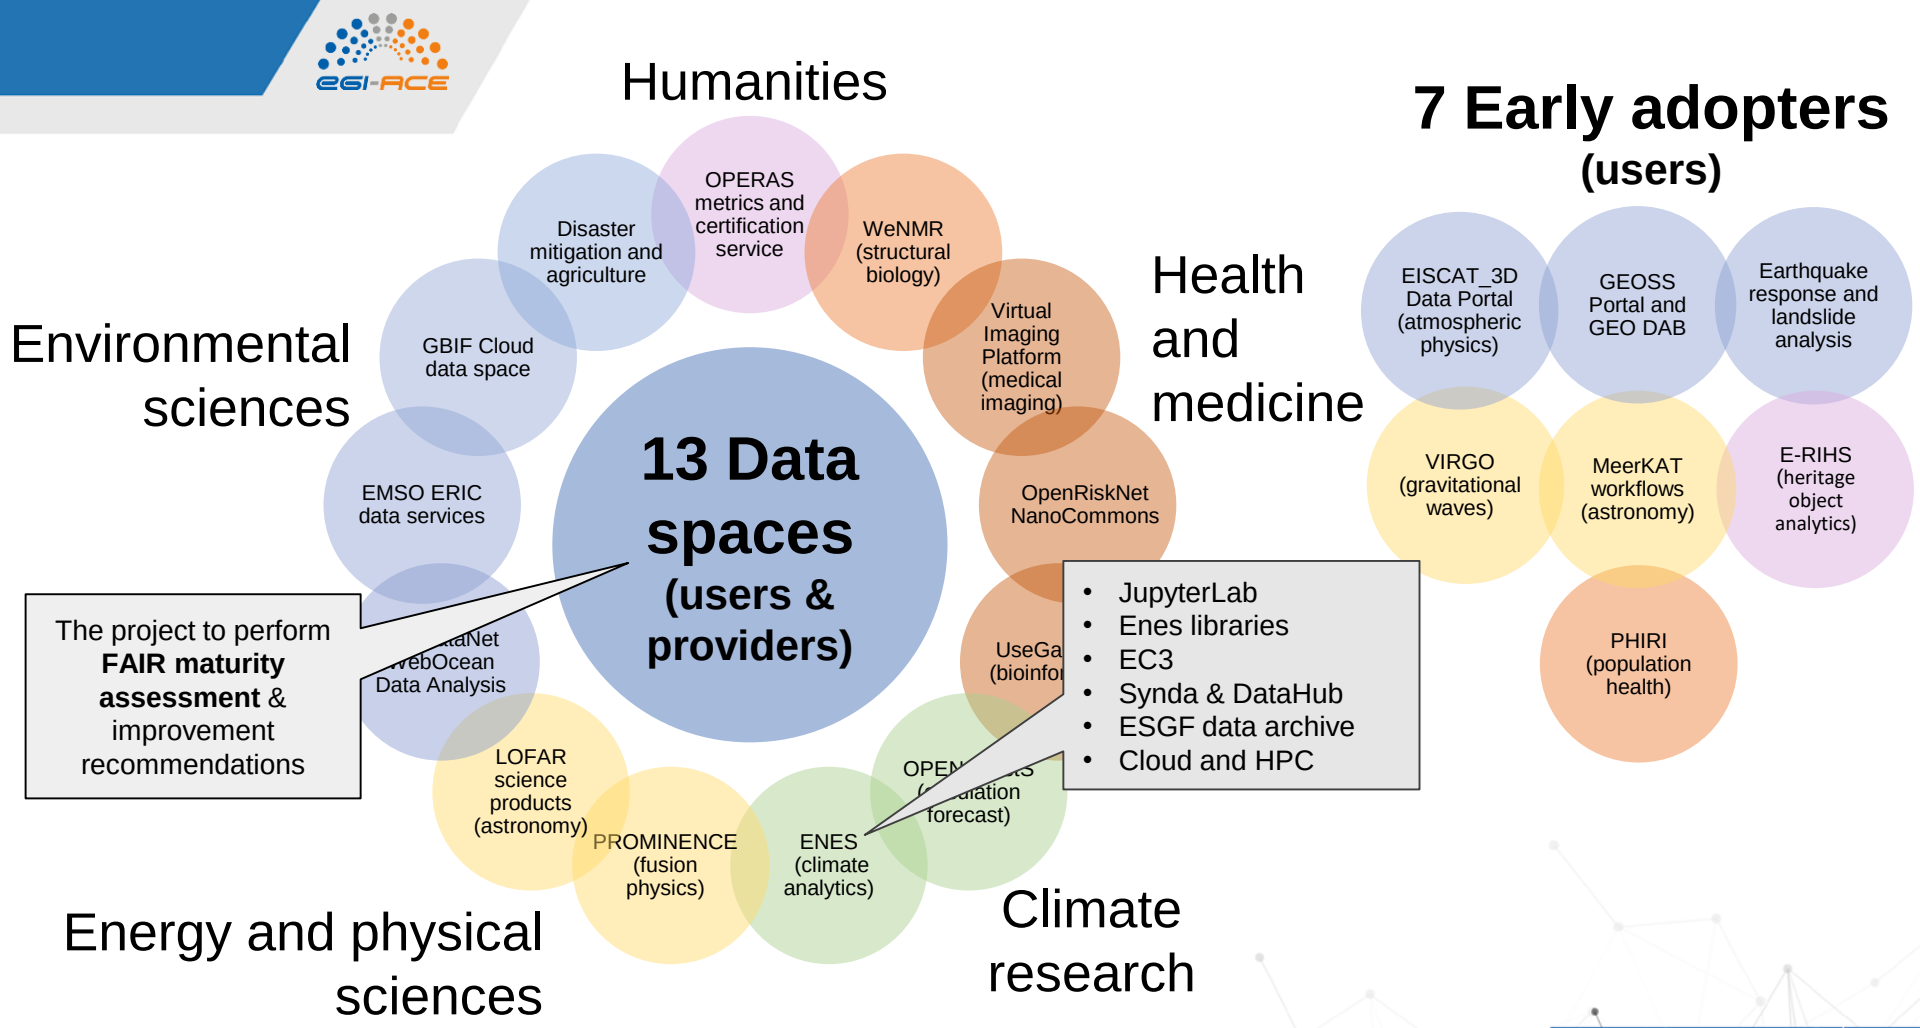
Humanities
7 Early adopters
(users)
Environmental sciences
Health and medicine
Energy and physical sciences
Climate research
13 Data spaces
(users & providers)
OPERAS metrics and certification service
Disaster mitigation and agriculture
WeNMR (structural biology)
GBIF Cloud data space
Virtual Imaging Platform (medical imaging)
EMSO ERIC data services
OpenRiskNet NanoCommons
EISCAT_3D Data Portal (atmospheric physics)
GEOSS Portal and GEO DAB
Earthquake response and landslide analysis
VIRGO (gravitational waves)
MeerKAT workflows (astronomy)
E-RIHS (heritage object analytics)
PHIRI (population health)
LOFAR science products (astronomy)
PROMINENCE (fusion physics)
ENES (climate analytics)
The project to perform FAIR maturity assessment & improvement recommendations
JupyterLab
Enes libraries
EC3
Synda & DataHub
ESGF data archive
Cloud and HPC

Possible interactions with INFRAEOSC07 projects

1. Close the research virtuous cycle
   - Interface EGI-ACE with long-term storage services (DICE, Zenodo, science spec.)
   - Interface EGI-ACE with publication services (Nexus - which one?)
2. Tracking impact (= interface with EOSC-Core?)
   - Extend/improve EGI Community Dashboard Beta to EGI-ACE (Nexus)
3. Connect Earth Observation services/data and EGI-ACE (C-SCALE, RELIANCE)
4. Coordinate user support, support ‘big communities’ together
   - FAIR-ness assessment of EGI-ACE Data Spaces (Who can help us?)
5. Green computing
   - Align practices
6. Communication
   - Cross-project promotion and dissemination

Summary

- **EGI-ACE delivers the Compute platform in EOSC**
  - Compute continuum with HTC, Cloud and HPC
  - Online data analytics services
- **We are open for business !**
  - Most of our services are already available in the EOSC Portal: https://marketplace.eosc-portal.eu/services?related_platforms=52
  - Call for use cases with 1st cut-off date to be published soon
- **Specific collaboration with OpenAIRE-Nexus to start next week**

Thank you

https://www.egi.eu/projects/egi-ace/

www.egi.eu
@EGI_eInfra

Infrastructure services: highlights from the compute, security, data areas

Andrea Manzi (EGI.eu)
Diego Scardaci (EGI.eu)
Enol Fernandez (EGI.eu)

The work of the EGI Foundation
is partly funded by the European Commission
under H2020 Framework Programme

Compute Area










Enol Fernández

EGI Cloud Compute

Federated IaaS

Distributed Infrastructure as a Service (IaaS) powered by the EGI Federated Cloud

- Execution of VM based workloads on a distributed infrastructure
- Federated identity
- Common VM image catalogue
- GUI and CLI/API based access
- Central accounting and monitoring

EGI Cloud Compute allows users to **host data- and compute- intensive applications** in a distributed infrastructure (move compute near the data) with **complete control** on the software and resources used

Expanding EGI Cloud with EGI-ACE

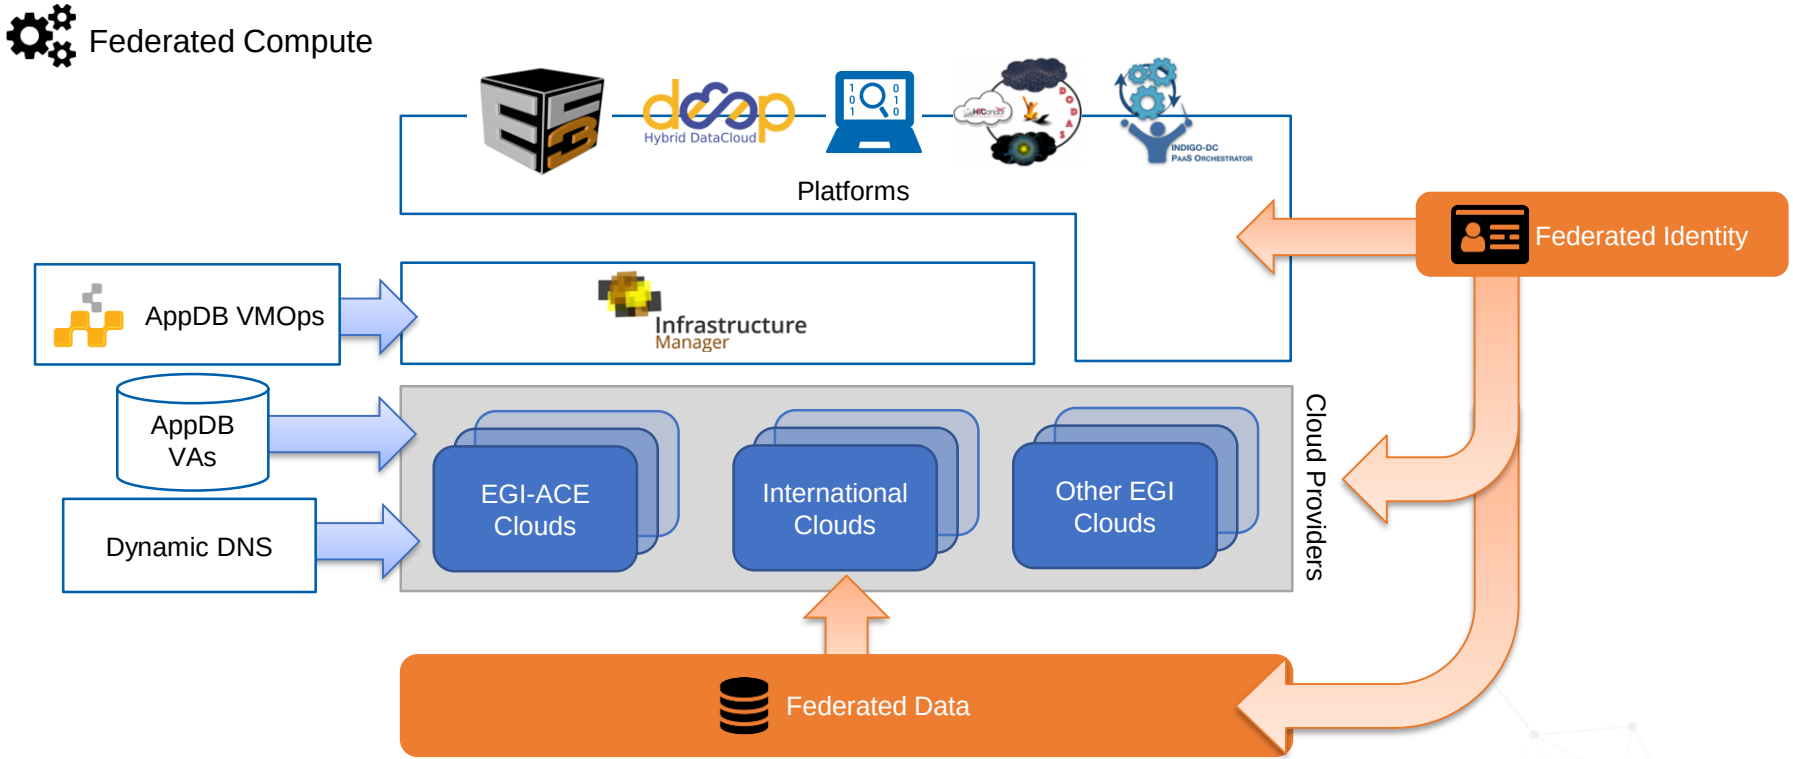

Bringing capacity for EOSC users

EC funded capacity (Virtual Access):

- **Phase 1 (month 1-10)** ⇒ 13.8 M CPU hours - **41,300 GPU hours** - 7.5 PB/month
- **Phase 2 (month 11-20)** ⇒ 27.6 M CPU hours - **83,000 GPU hours** - 15 PB/month
- **Phase 3 (month 21-30)** ⇒ 41.4 M CPU hours - **124,000 GPU hours** - 23 PB/month

\+ Policy-based and pay-per-use access from members of EGI Federation, international partners and commercial clouds

Orchestration

IM deploys virtual infrastructures on Cloud resources

- Infrastructure as Code (IaC) / RADL & TOSCA descriptions
- Automates deployment, configuration, software installation, monitoring and update of virtual infrastructures
- Wide variety of back-ends thus user applications become cloud agnostic

Provisioning of virtualized compute and storage on distributed infrastructure: cloud, HPC and container orchestration platforms

- Automatic selection of providers
- Federated Data integration: data orchestration, replication, trigger jobs on events
- Automatic re-submission of deployments

Automated Cluster Management

Elastic Cloud Computing Cluster (EC3) deploys virtual **elastic** clusters on IaaS providers

- IM for deployment, CLUES for elasticity
- Predefined ansible templates: Kubernetes, Mesos, SLURM, Torque, SGE, HTCondor, Nomad...

DODAS supports the deployment and configuration of distributed (possibly federated) clusters for executing experiment workflows (e.g. data processing, data analysis):

- Batch System on demand
- Big Data clusters on demand Interactive analysis service for local and remote data processing

AI/ML: DEEP training facility

1-click development and training environment in cloud and HPC resources

- Train, test and evaluate of models through the DEEPaaS API
- Transparent access to infrastructure resources
- Based on TOSCA templates submitted via INDIGO-DataCloud PaaS orchestrator and Infrastructure Manager

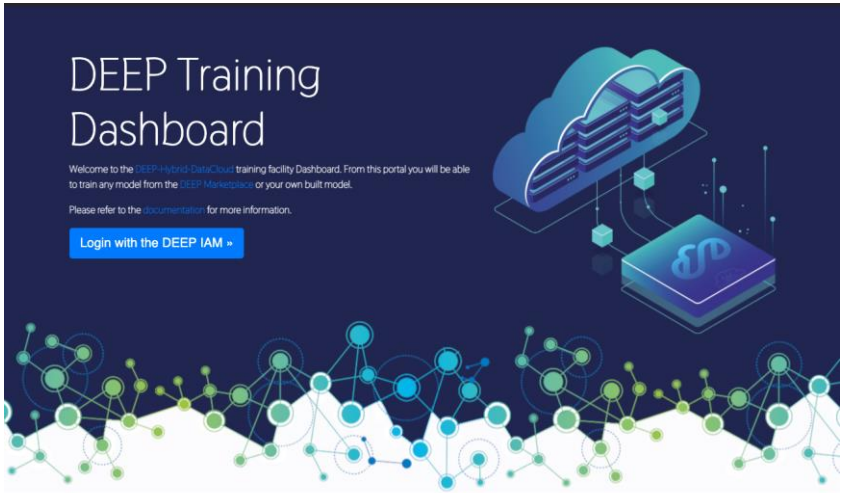

Interactive Computing: Notebooks

JupyterHub hosted in the EGI Cloud

- Jupyter notebooks ‘as Service’
- One-click solution: login and start using

EGI-ACE will bring:

- Binder (reproducible notebooks) support
- Better integration with Federated Data (e.g. onedata) and Federated Compute (DIRAC) services

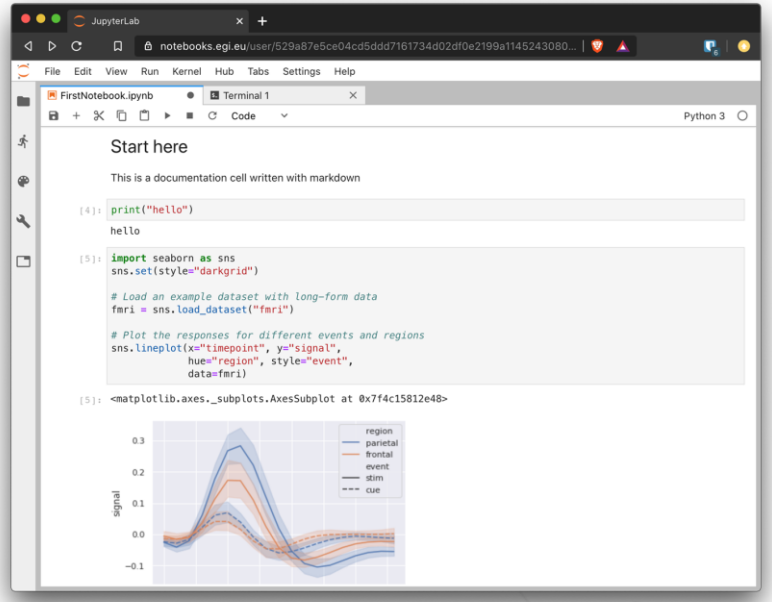

Beyond IaaS Clouds: HTC and HPC

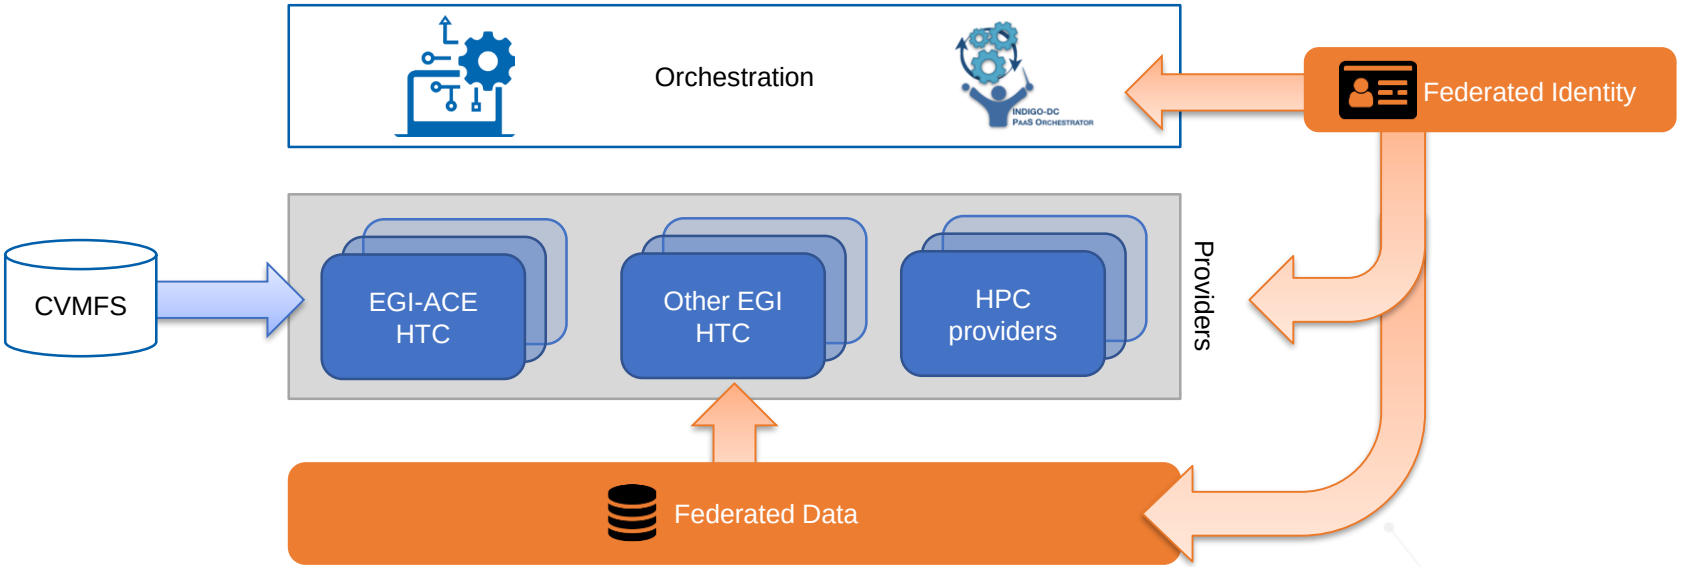

High Throughput Compute

The EGI High-Throughput compute (HTC) provides users with the capability to access large amounts of computing resources, and to submit hundreds or thousands of computational tasks.

EGI-ACE brings one new option to the HTC service from SURF:

- Spider, a versatile high-throughput data-processing platform aimed at processing large structured data sets.

Workload Manager

The Workload Manager helps users to exploit distributed resources in a transparent way.

- Based on DIRAC, a framework shared by multiple experiments, both inside HEP, astronomy, and life sciences
- Optimises workload placement
  - Support for massive workload and data operations
  - Aggregates multiple types of resources, e.g. Cloud / HTC providers, user community resources
- User-friendly GUI, CLI and API access for advanced usage
  - Jupyter notebooks (prototype)
- Services customizable for the needs of particular communities

Software Distribution

CernVM-File System (CVMFS) provides a POSIX read-only file system in user space

- Scalable software distribution, recently extended to handle container images
- Caching at providers for fast access to frequently used software
- Mostly used in HTC, but useful also for IaaS

Application Database (AppDB) provides Virtual Appliance (VAs) Catalogue

- Public registry of images with detailed metadata
- Community-managed VA lists, fine control on what can be executed by users
- Automatic distribution as VM images to the cloud providers

OPENCoastS: forecast systems for coastal sites

OPENCoastS brings together several EGI services to deliver on-demand coastal circulation forecast systems through a web platform with minimal user intervention

- Hosting of web frontends / databases and permanent services on Cloud Compute
- Dynamic workload managed by Workload Manager and executed on High Throughput Compute

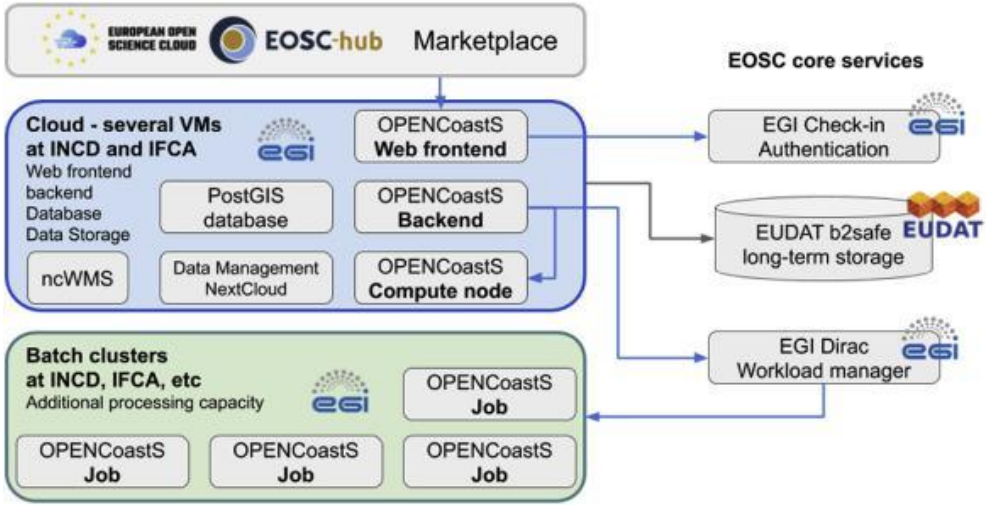

A distributed virtual CMS site powered by DODAS

DODAS used to build a lightweight ephemeral WLCG-Tier on demand distributed across four Cloud providers

- Executing actual experiment workload
- Dynamically distributing the load as providers join
- Uses PaaS Orchestrator to deploy elastic Kubernetes that run HTCondor + CVMFS configured for integrating into CMS submission infrastructure

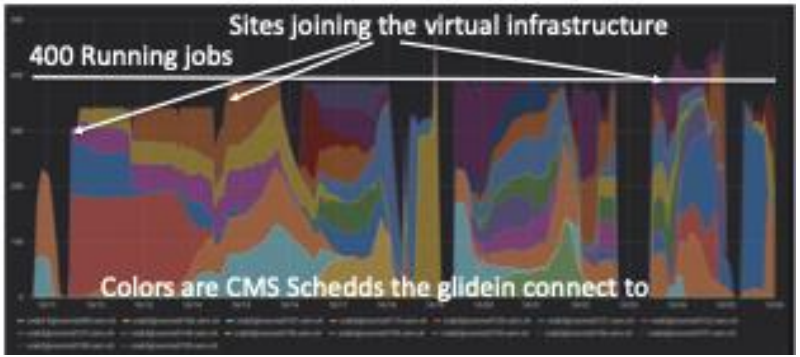

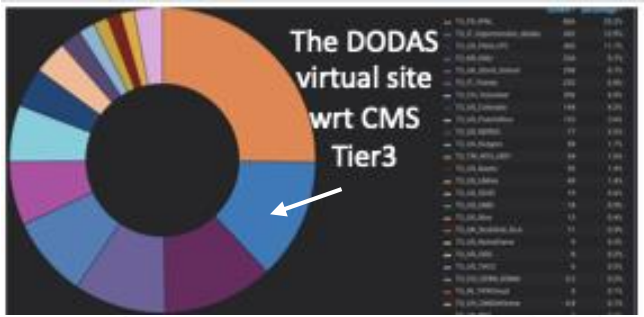

Security Area

Check-in

RCAuth.eu &
Master Portal

Diego Scardaci

EGI Check-in

Proxy service that operates as a central hub to connect federated Identity Providers (IdPs) with EGI service providers

Perun

RCAuth.eu & Master Portal

Main characteristics:

- Enables multiple federated authentication sources using different technologies
- Federated in eduGAIN as a service provider (REFEDS RnS and Sirtfi compliance)
- User registration portal to allow accounts-linking
- Combines user attributes originating from various authoritative sources (IdPs and attribute provider services)
  - Integrated with **PERUN**
- Compliant with the AARC Blueprint
- Support non-web use cases & delegated access

The EGI AAI Platform: Check-in

The IdP/SP Proxy

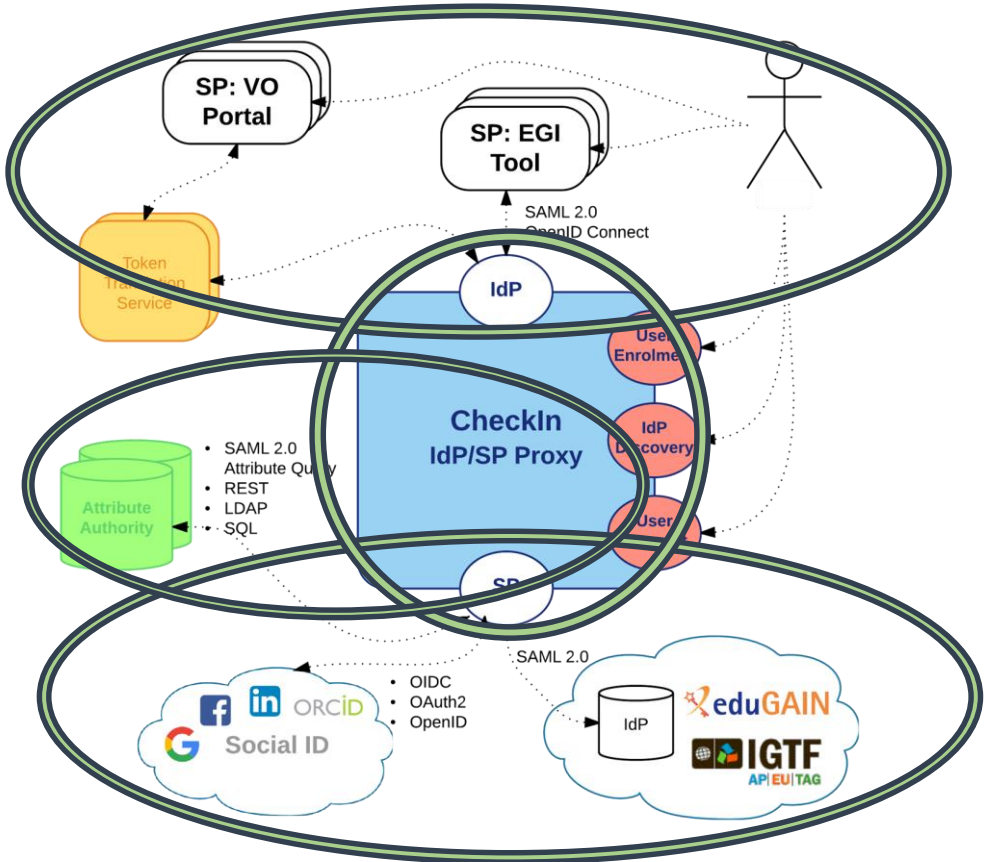

- Identity Providers:
  - SAML2.0: eduGAIN
  - OIDC/OAuth2: Google, Facebook, LinkedIn, ORCID, WeChat, BitBucket
  - X.509: IGTF
- Service Providers:
  - SAML2.0 & OIDC
- Attribute Authorities
  - SAML2.0 Attr. Query, REST, LDAP, SQL
- Token Translation Services
  - SAML2.0-to-X.509: Master Portal to RCauth.eu Online CA
- Support for Levels of Assurance
- User enrolment & account linking
- IdP Discovery
- User Consent

Use case: AAI as a Service

Check-in as an authentication proxy to allow user logins from institutional IdPs in eduGAIN and social media for non EGI services

- Minimal overhead for the service development
- Prerequisites:
  - Service provider must accept EGI policies on data protection

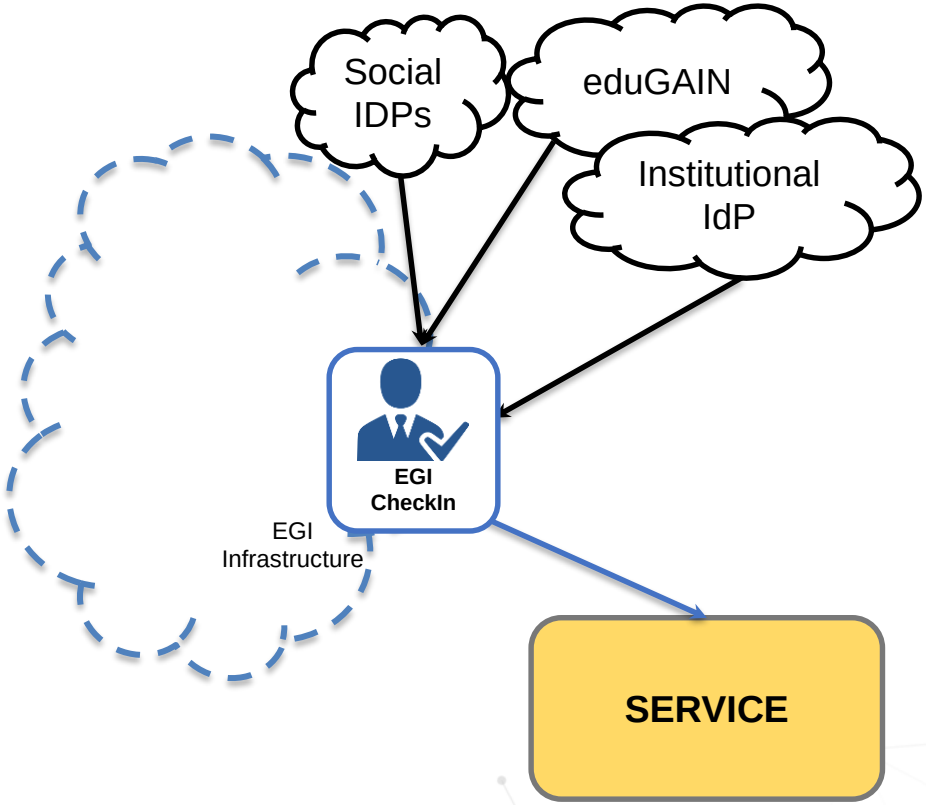

Use case: AAI Integration

Community operating its own AAI connected as an IdP to Check-in to allow its users to access EGI services & resources

- Users can access EGI services without changing their authentication workflow

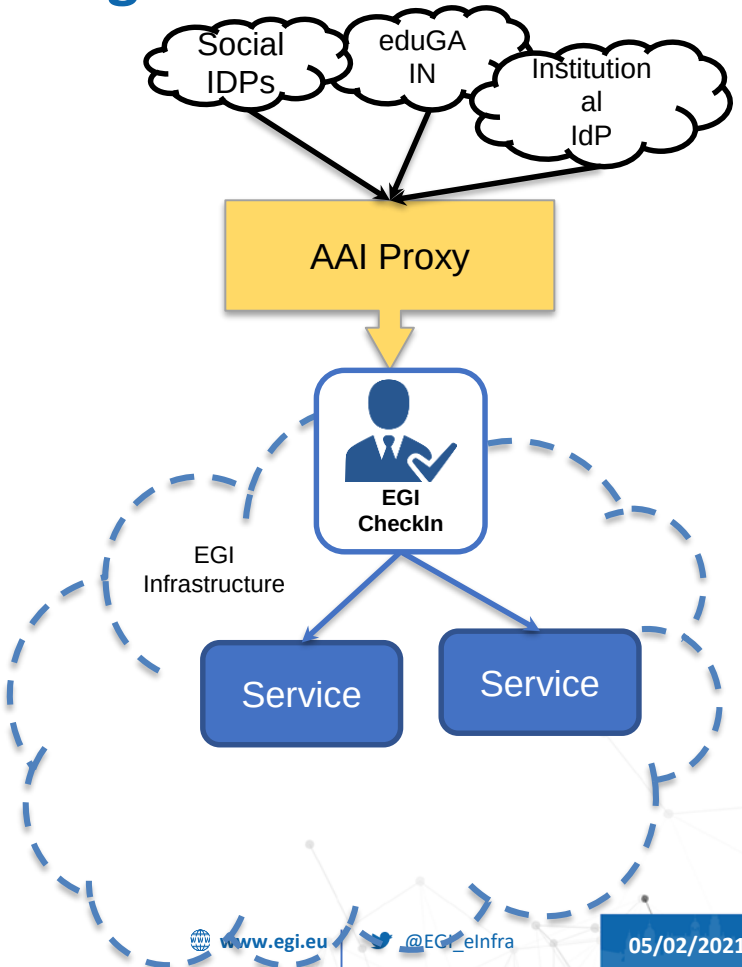

Use case: External Attribute Provider

Community managing authorisation information about the users (VO/group memberships and roles) via their own group management service, which is connected to CheckIn as an external attribute authority

- CheckIn will handle the configuration of the IdPs and the aggregation of the attributes for the SPs

- No need to migrate the group information into an EGI specific attribute authority

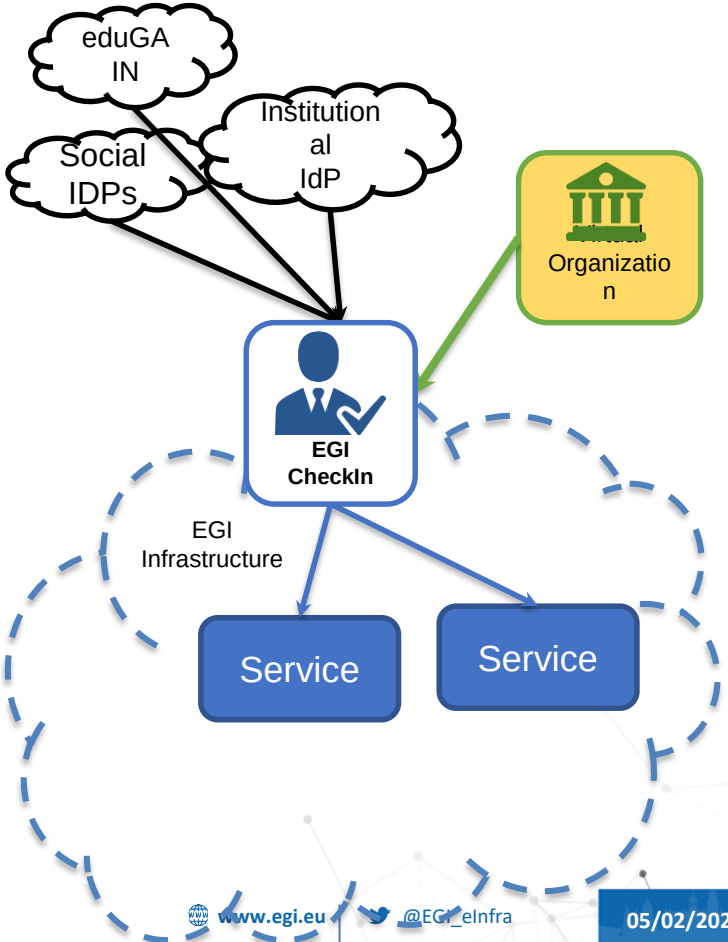

Use Case Group Management as a Service

Communities that do not operate their own group management service can leverage the group management capabilities of the CheckIn platform

- Avoid overhead of deploying a dedicated group management service
- Authorised VO admins will manage the information about their users independently
- Can be used with EGI and non-EGI services

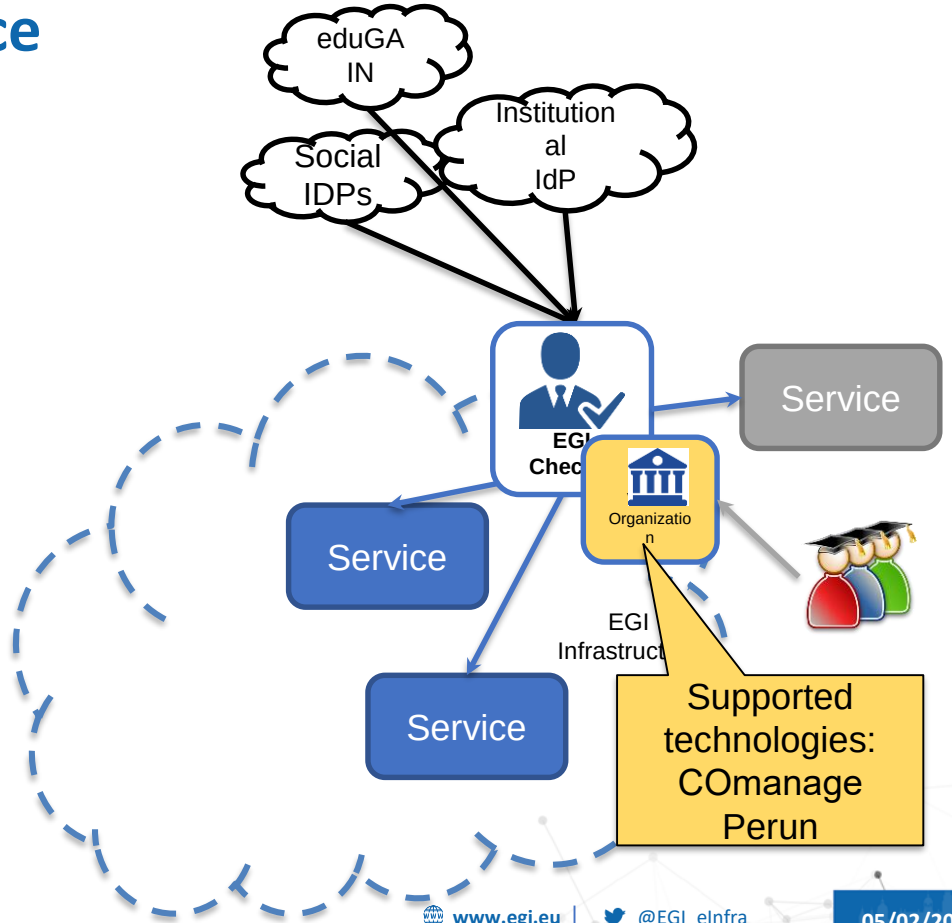

EGI Check-in Statistic viewer

Login statistics for services integrated

EGI Check-in statistics

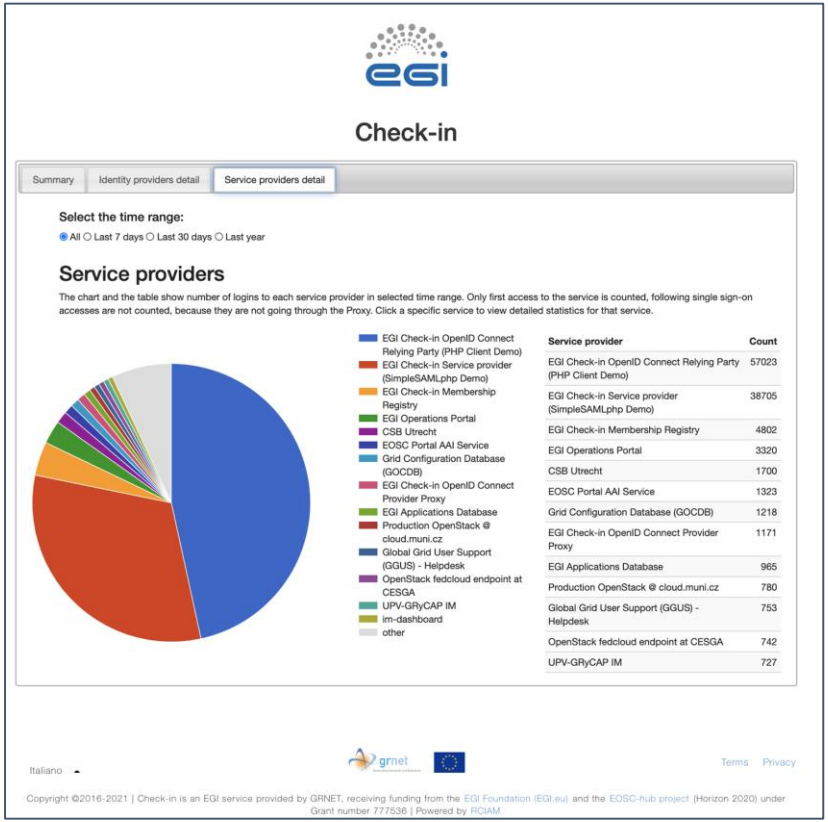

General view of the login statistics (all services)

Dedicated service view for login statistics:

- time trend of user logins
- distribution of logins based on the IdPs used by users

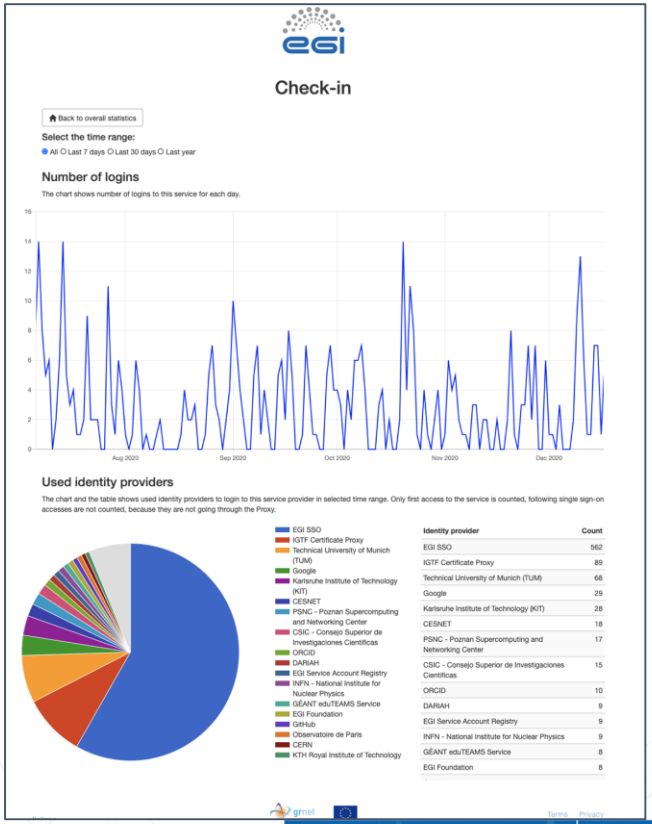

Some examples: EISCAT-3D and OPERAS

EISCAT-3D

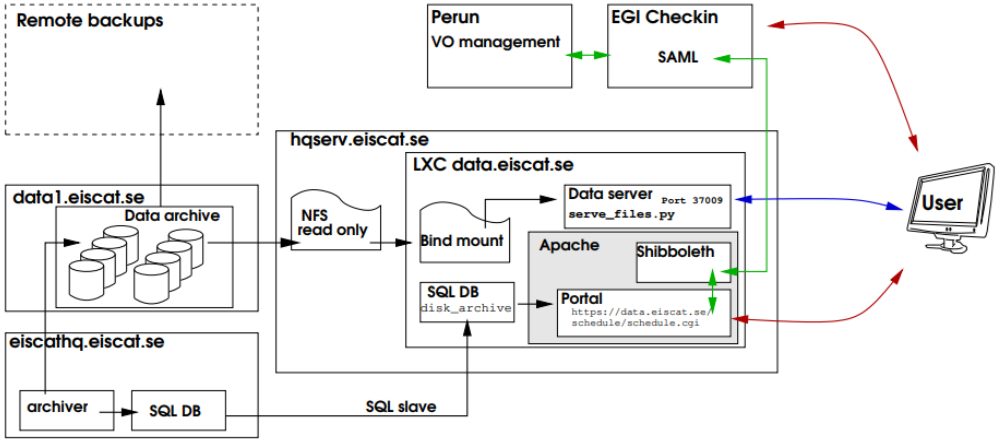

Radars for the upper atmosphere and ionosphere

- EISCAT-3D users managed in Perun
- Access to EISCAT services through Check-in

OPERAS

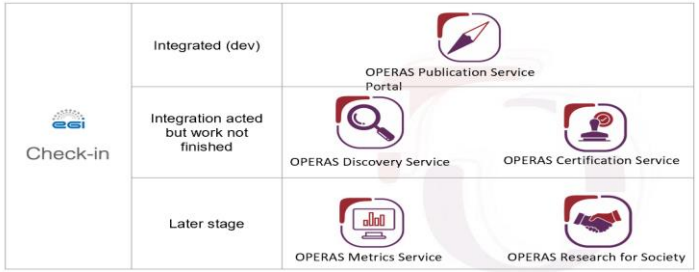

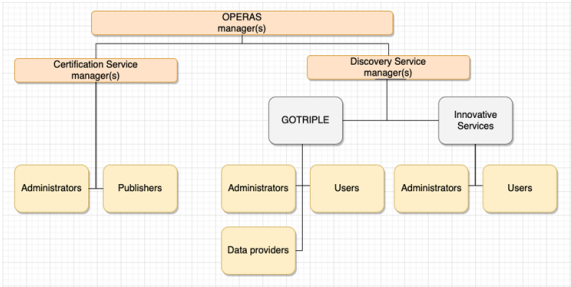

Open Scholarly Communication in the ERA for Social Sciences and Humanities

- OPERAS users managed in Comanage and structured in groups and roles

Check-in Documentation

- Usage guide
- Integration guide for service providers
- Integration guide for identity providers
- Getting Support

Data Area




Andrea Manzi

Online Storage

Includes 3 categories of storage services

- File Storage
- Block Storage
- Object Storage

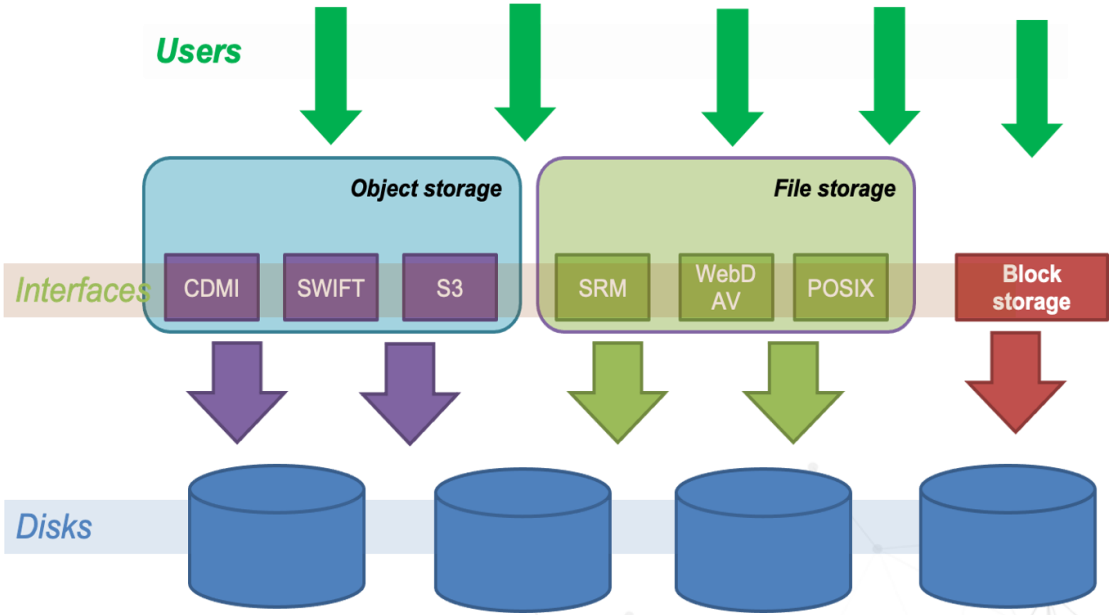

Online Storage

File storage

- File storage can be used for storing and accessing files on the infrastructure as input/output to EGI HTC computations
- The EGI Workload Management system (DIRAC) is able to access files stored in File Storage instances via different protocols, and schedule the computations in order to be executed close to where the input files are stored
- Technologies: DPM, dCache, StoRM

- Interfaces: SRM, HTTP/WebDAV, XRootD, gridftp, CDMI

Online Storage

Block storage

- Block-level storage solution that allows to **expand** the storage capacity of instances in the EGI Federated Cloud, offering the lowest possible latency for applications
- **Increase** storage without increasing the size or capacity of the instance or by provisioning new ones; delete servers, **keeping data intact**
- Enhancing Technology: Openstack Cinder
- Interfaces: POSIX for access, OCCI and Openstack CLI/GUI for management

Online Storage

Object storage

- Manages data as **objects**. Each object includes the data itself, a variable amount of metadata, and a globally unique identifier.
- Cloud object storage allows relatively inexpensive, scalable and self-healing retention of massive amounts of unstructured data
- the main uses cases are:
  - Backups and Data Archive, AI and ML , Cloud Native applications
- Technologies: OpenStack Swift, CEPH
- Interfaces: S3, Swift, CDMI

DataHub

Access key scientific datasets in a scalable way

- Service based on **Onedata** technology

- It allows **transparent data access** under a **common namespace** regardless of the location
- Data can be accessed via a **GUI** or **APIs**
- Developed in EOSC-hub project and under integration by many communities
- Central component and distributed

providers in EGI-ACE are going to be operated by **Cyfronet**

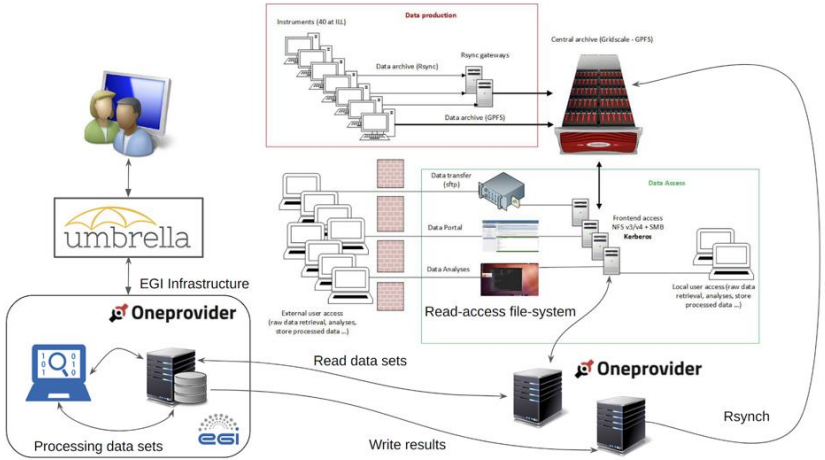

DataHub

Main functionalities

- Integrated with EGI-Checkin
- File/Share Management
  - Shares can eventually be published via Handle services and have a PID assigned
- Extensive backend support
  - Posix, S3, Swift, CEPH, GlusterFS
- Data import
  - Allows exposing existing datasets
- Metadata Management
  - Metadata associated to files (k-v pairs, json, rdf)
  - OAI-PMH interface
- Harvester Management
  - Possibility to automatic extract metadata from files published to a space
  - Metadata updates are pushed to external indexes (e.g. Elastic Search) and automatically synchronized

Data Transfer

Transfer large datasets between storages

- Easy users interaction for submitting transfers
- Portal for end-users, Real Time monitoring and Web Admin
- Checksums and retries are provided per transfer

- Multiprotocol support
- Different clients to access the service (RESTFul APIs, python bindings)

- Transfers from/to different storages (including staging from tape)
- Parallel transfers scheduling and optimization to get the most from network without burning the storages

- Priorities/Activities support for transfers classification

Data Transfer

Service operations

- Orchestrates transfers between EGI Online storage instances and external storages using standard protocols (e.g. HTTP/WebDAV, gridftp)

- **STFC-UKRI** is providing the service and planning to operate the end user graphical interface (WebFTS)

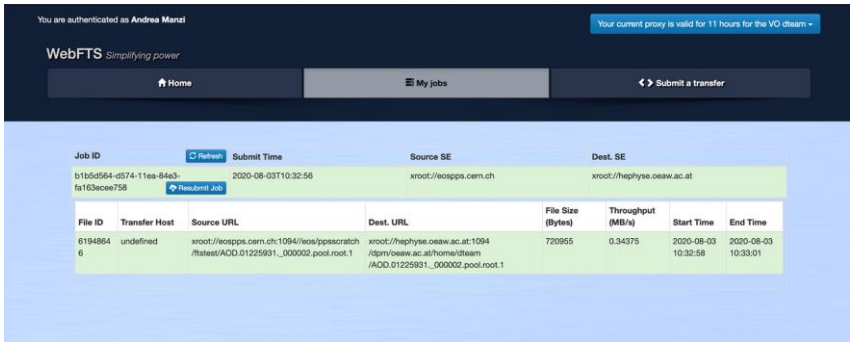

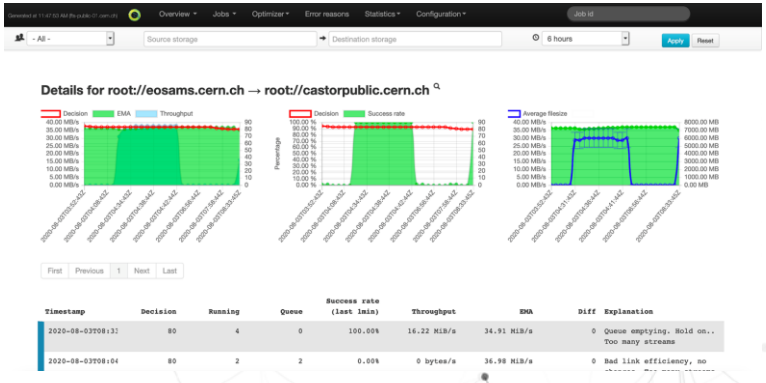

Rucio

Scalable data management

Rucio is the software developed at CERN for the management of LHC experiments data

- provides a complete and generic scientific data management service
- Rucio manages multi-location data in a heterogeneous distributed environment
  - Creation, location, transfer, and deletion of replicas of data
  - Orchestration according to both low-level and high-level driven data management policies (usage policies, access control, and data lifetime)
  - Interfaces with workflow management systems
  - Large-scale and repetitive operational tasks can be automated

Rucio

Service operations

- Under integration by many communities also outside HEP
  - Largest user (ATLAS) managing 500PB of data and 20GB/s transfers between 72 official sites

- **STFC-UKRI** implemented **extensions for Multi-VO support** and will operate the service

- Interfaced to FTS for transfer execution and possibly integrated with EGI Workload manager to implement data transfer orchestration linked to workflow executions

OpenRDM

Main functionalities

A **complete solution for managing research data**, based on the openBIS software.

Targeting research labs working with quantitative research data

- **Inventory** management
  - Keep track of all materials and samples used in a lab
- **Electronic lab** notebook
  - Describe computational experiments and links to the materials, samples and protocols stored in the inventory
- **Data analysis** integration
  - Jupyter notebook integration and access to external computing clusters
- **Data management**
  - Store all data connected to experiments, export to data repos like Zenodo

OpenRDM

Service Operations

- The service will be provided by **EnhanceR** as a centrally managed instance (**OpenRDM.eu)** and will provide access to **long-tail of science** type use cases

- Additionally dedicated instances for specific communities will be setup on top of EGI Federated resources

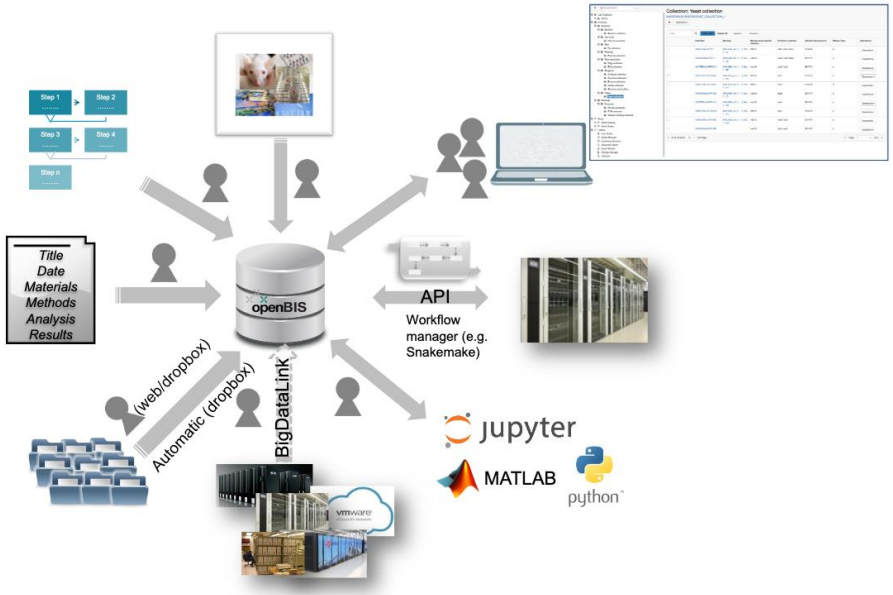

Documentation

Internal and External

- Online Storage
- DataHub (user and provider guides)
- Data Transfer
- Rucio
- OpenBIS user guide