EGI technical platforms for advanced computing

Tiziana Ferrari

Technical Director, EGI Foundation

www.egi.eu

EGI-Engage is co-funded by the Horizon 2020 Framework Programme of the European Union under grant number 654142

EGI-Engage

Outline

- Introduction to EGI
- Services for distributed computing, data management and AAI
- New requirements, new challenges
- Towards the European Open Science Cloud

The EGI Services are provided by the EGI Federation

EGI Council participants: national e-infrastructure providers and international research organisations (CERN and EMBL)

Integrated e-infrastructure providers

826,000
Cores of compute capacity

560,000
Terabytes of storage capacity

48,000
Users

15
Research Infrastructures integrated with EGI

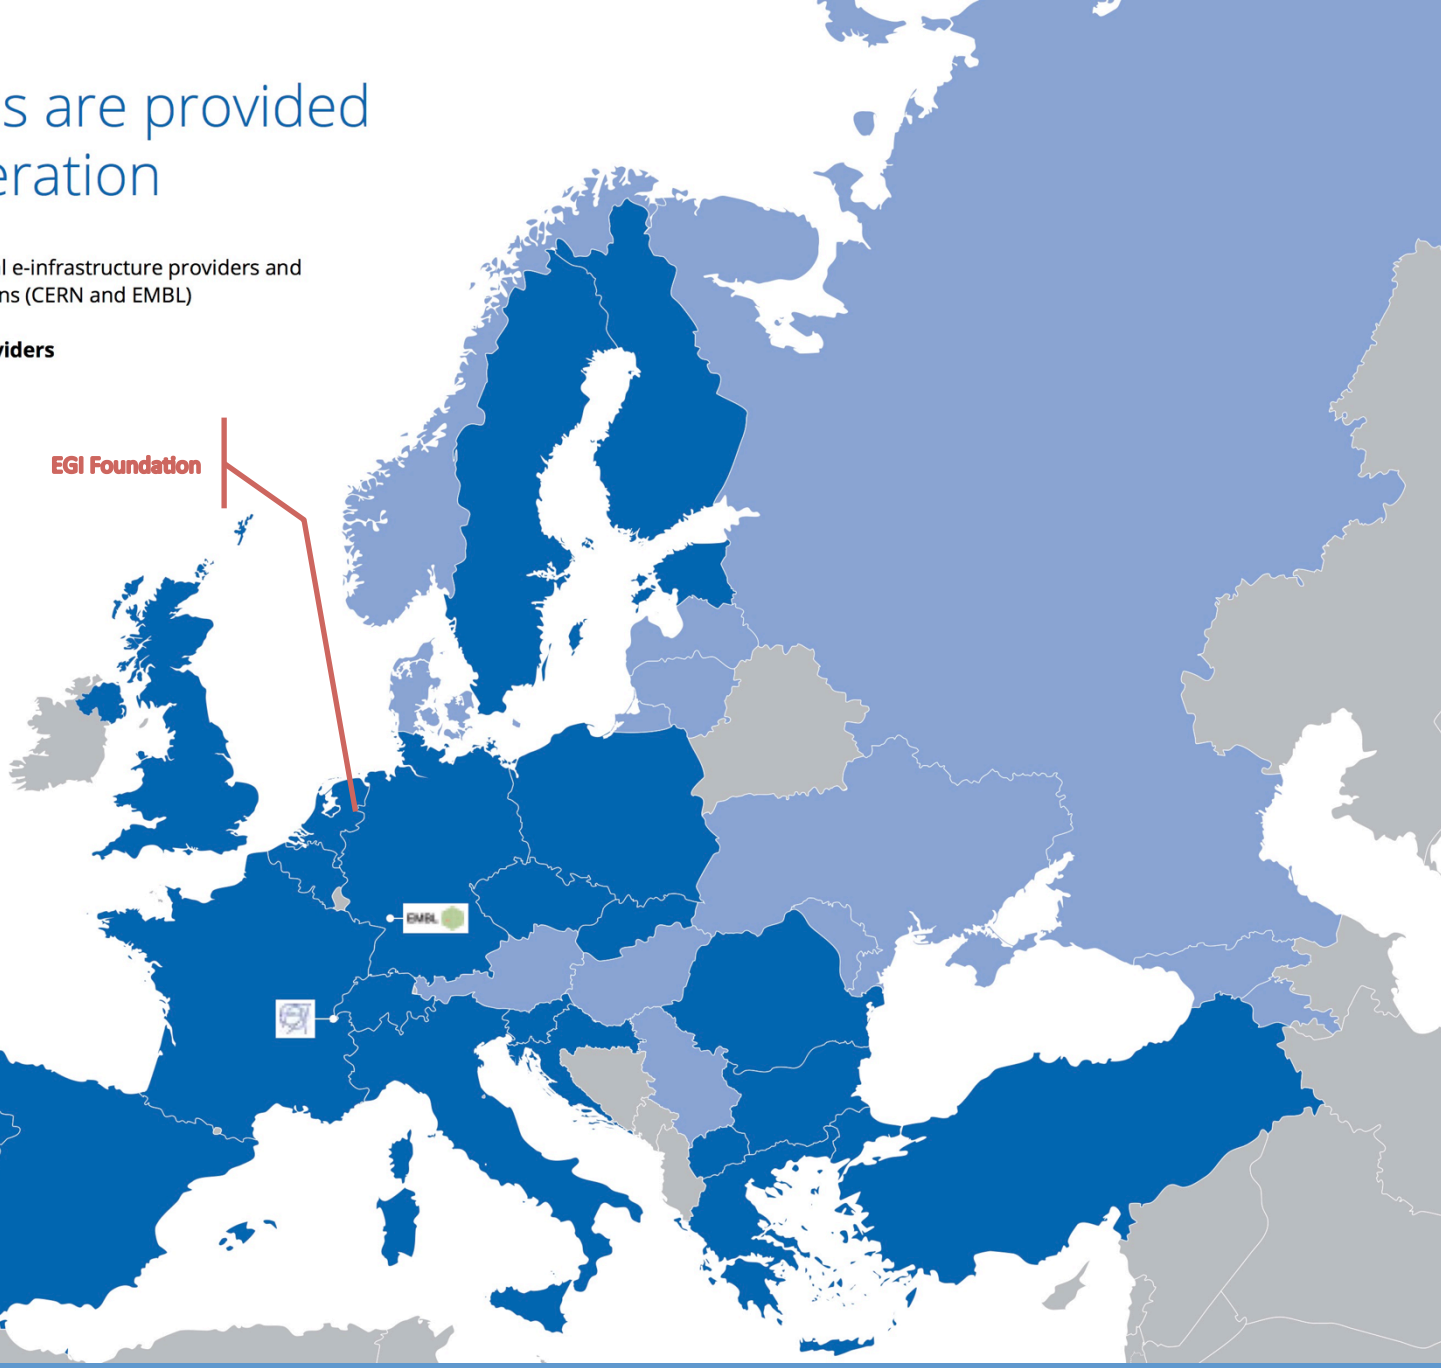

EGI Membership

https://www.egi.eu/about/egi-foundation/

- Major national e-Infrastructures: 22 NGIs
- EIROs: CERN and EMBL-EBI
- EGI Foundation
- (ERICs)

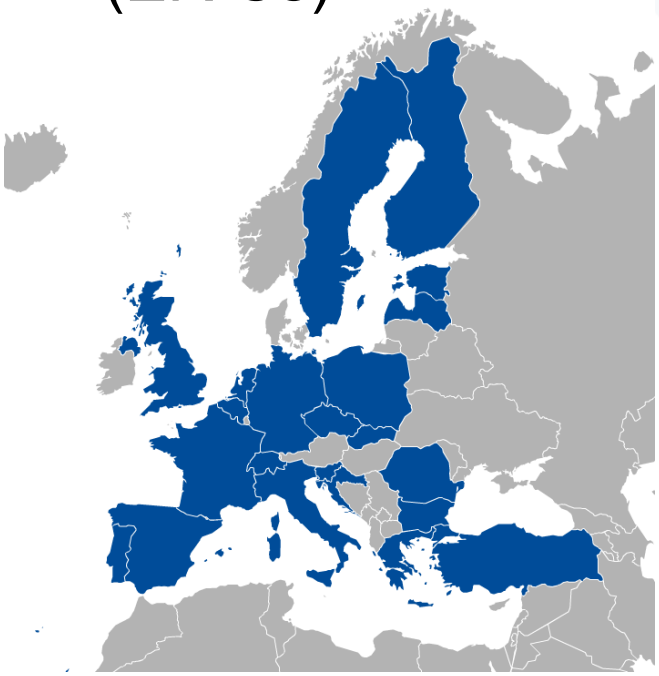

International Partnerships

Canada

USA

China Inst. Of HEP Chinese Academy of Sciences

India Centre for Development of Advanced Comp.

Africa and Arabia
Council for Scientific and Industrial Research, South Africa

Asia Pacific Region
Academia Sinica at Taiwan

Latin America
Universida de Federal do Rio de Janeiro

Ukraine
Ukrainian National Grid

EGI Federation, 2016 QR3

The largest distributed compute e-Infra worldwide

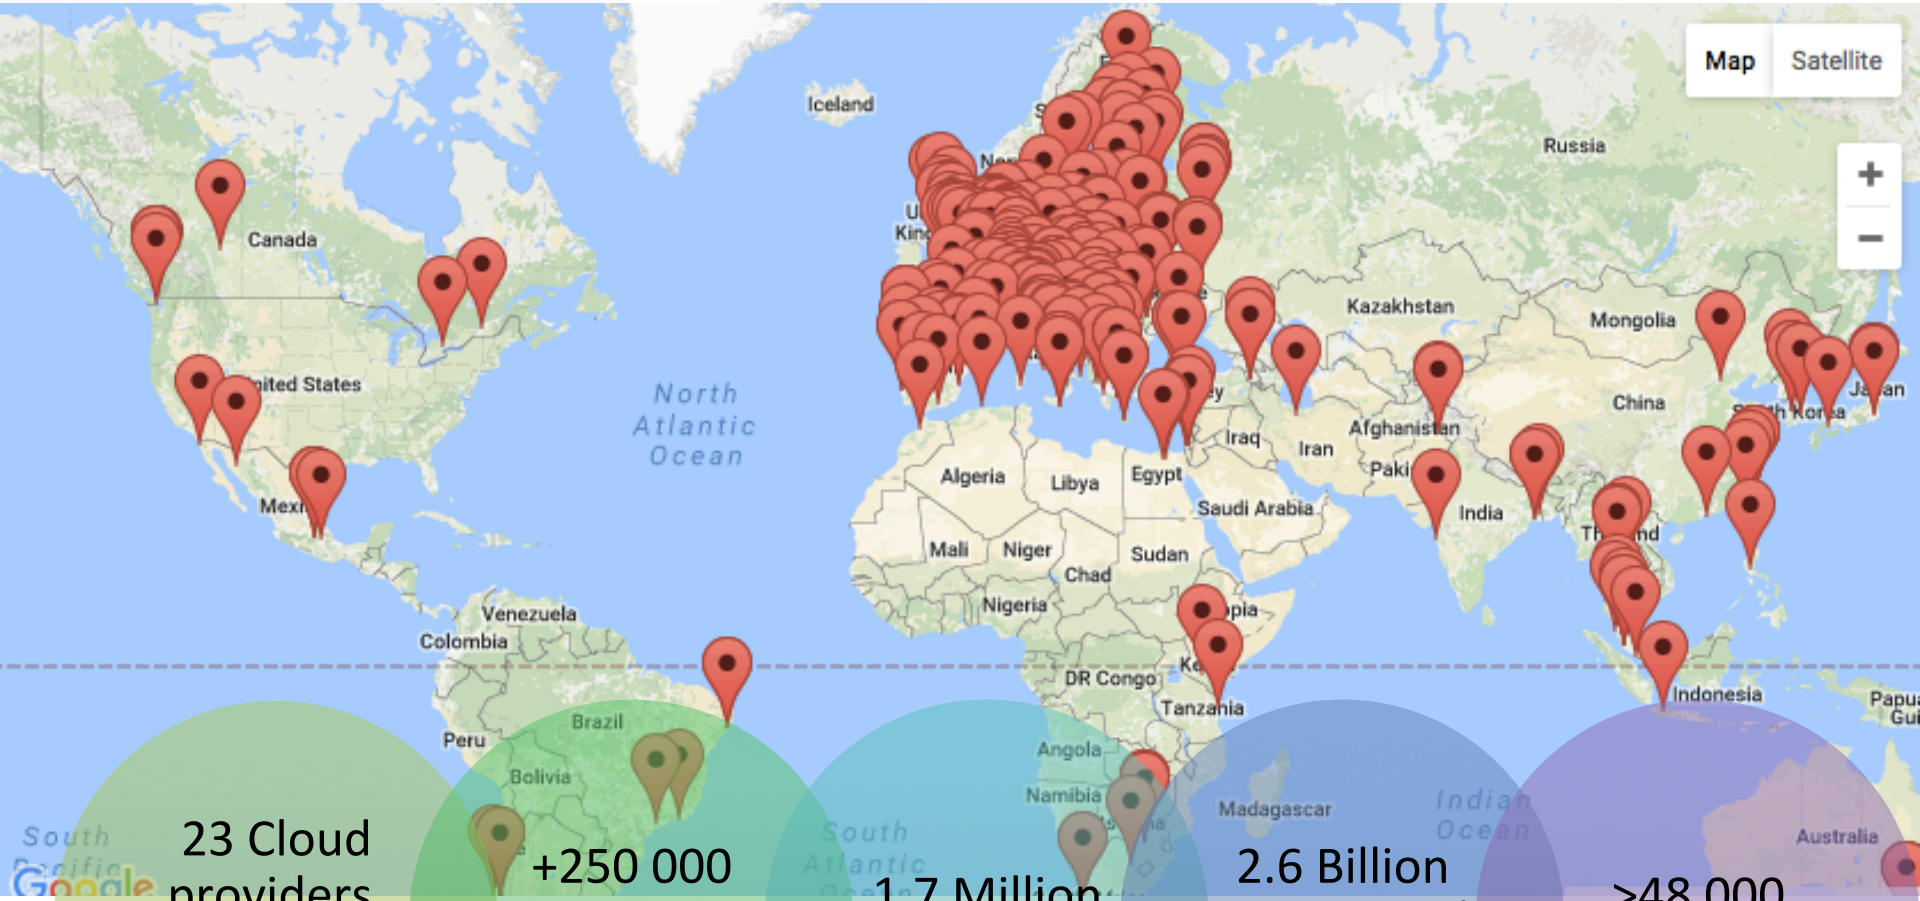

23 Cloud providers, +300 data centres

+250 000 instantiated VMs/year

1.7 Million jobs/day

2.6 Billion CPU hours/ year +26%

>48 000 users, +25%

Serving researchers and innovators

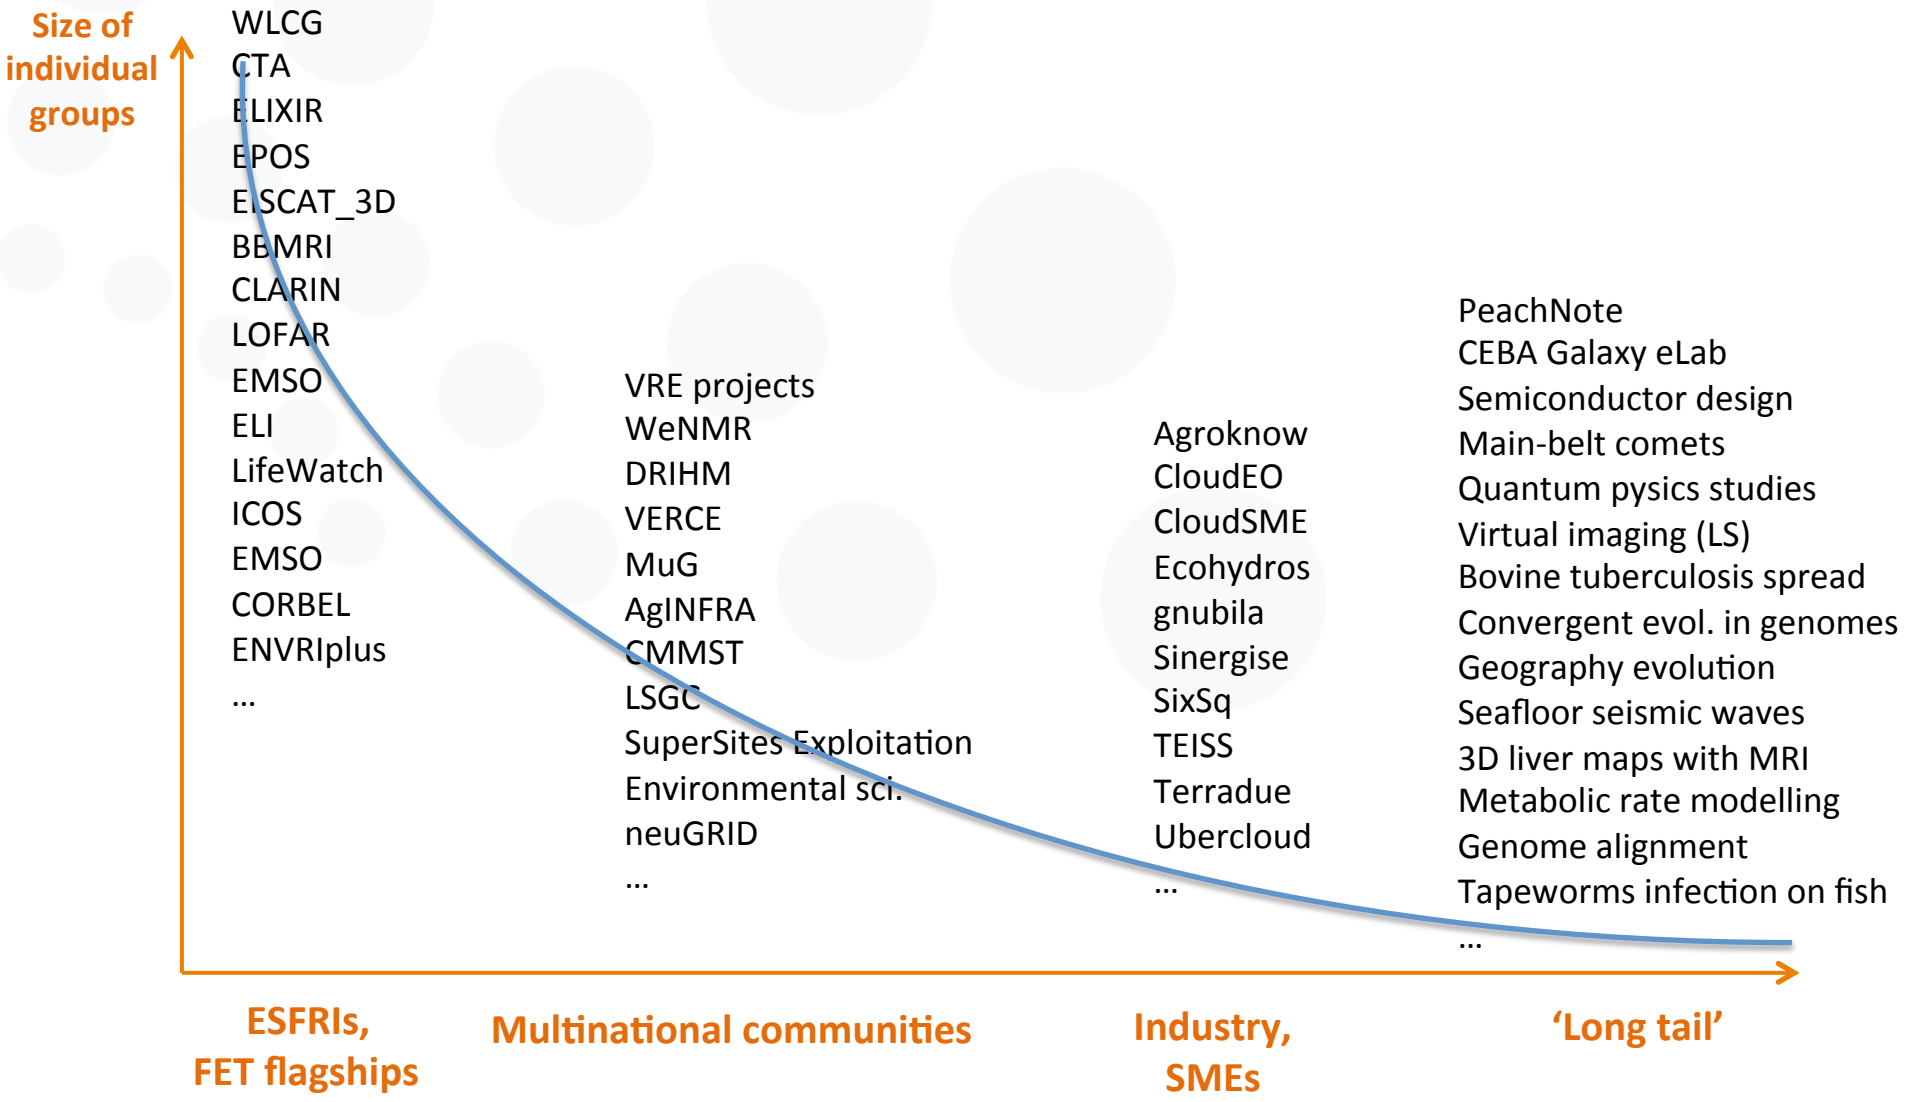

Supporting international research communities and thematic services

Installed compute capacity trends 2011-2016

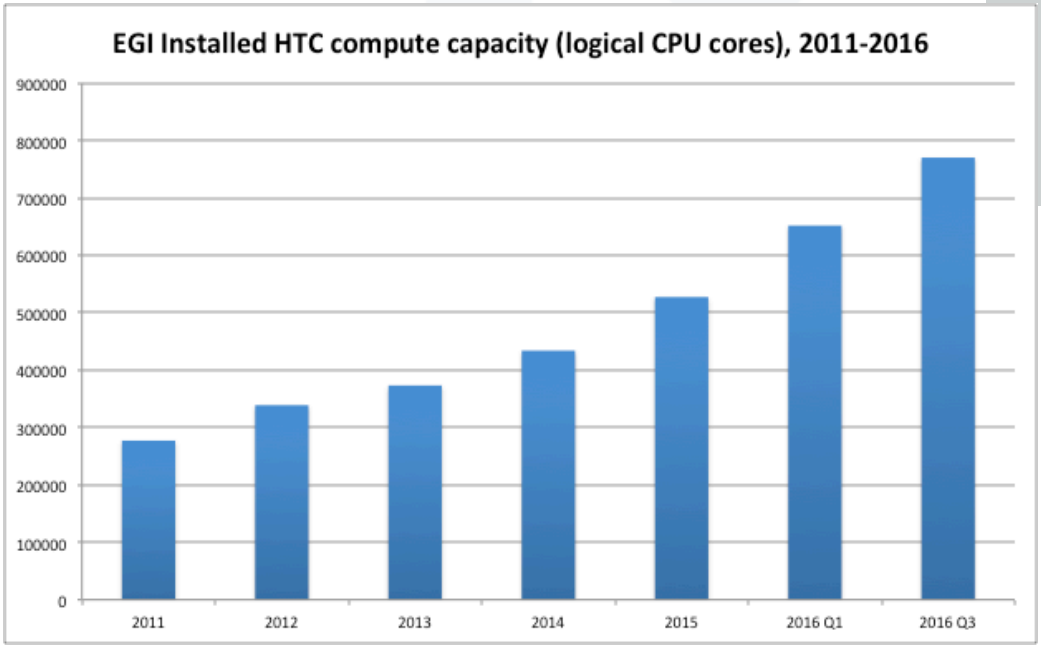

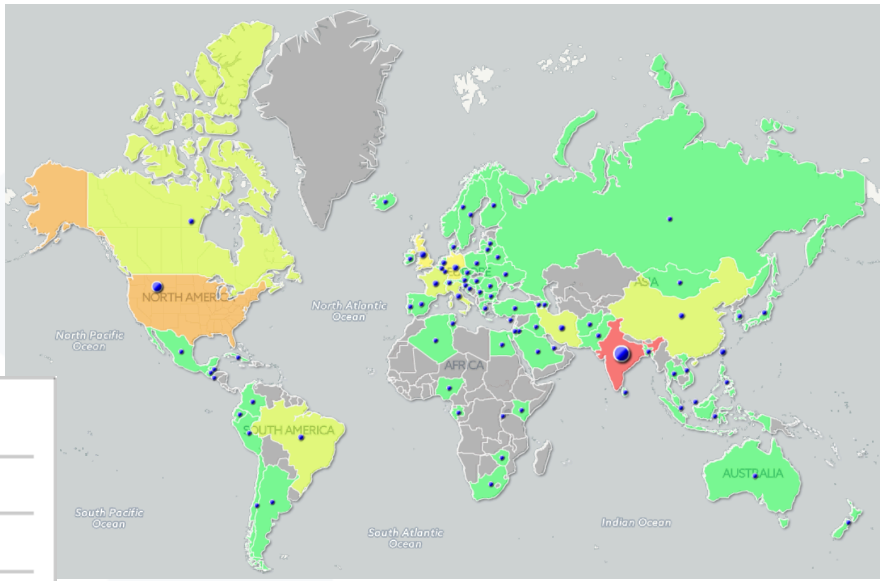

Example: Structural Biology
Distribution of users (2016, QR3)

- 2700 users
- 81 countries

(credits: A. Bonvin, WeNMR)

Astronomy/Astrophysics/Astro-particle physics projects and RIs in EGI

ARGO, AUGER, CTA, KM3NeT, LHCb, LOFAR, Large Synoptic Survey Telescope/LSST, PAMELA, ESA Planck Mission, XENON etc.

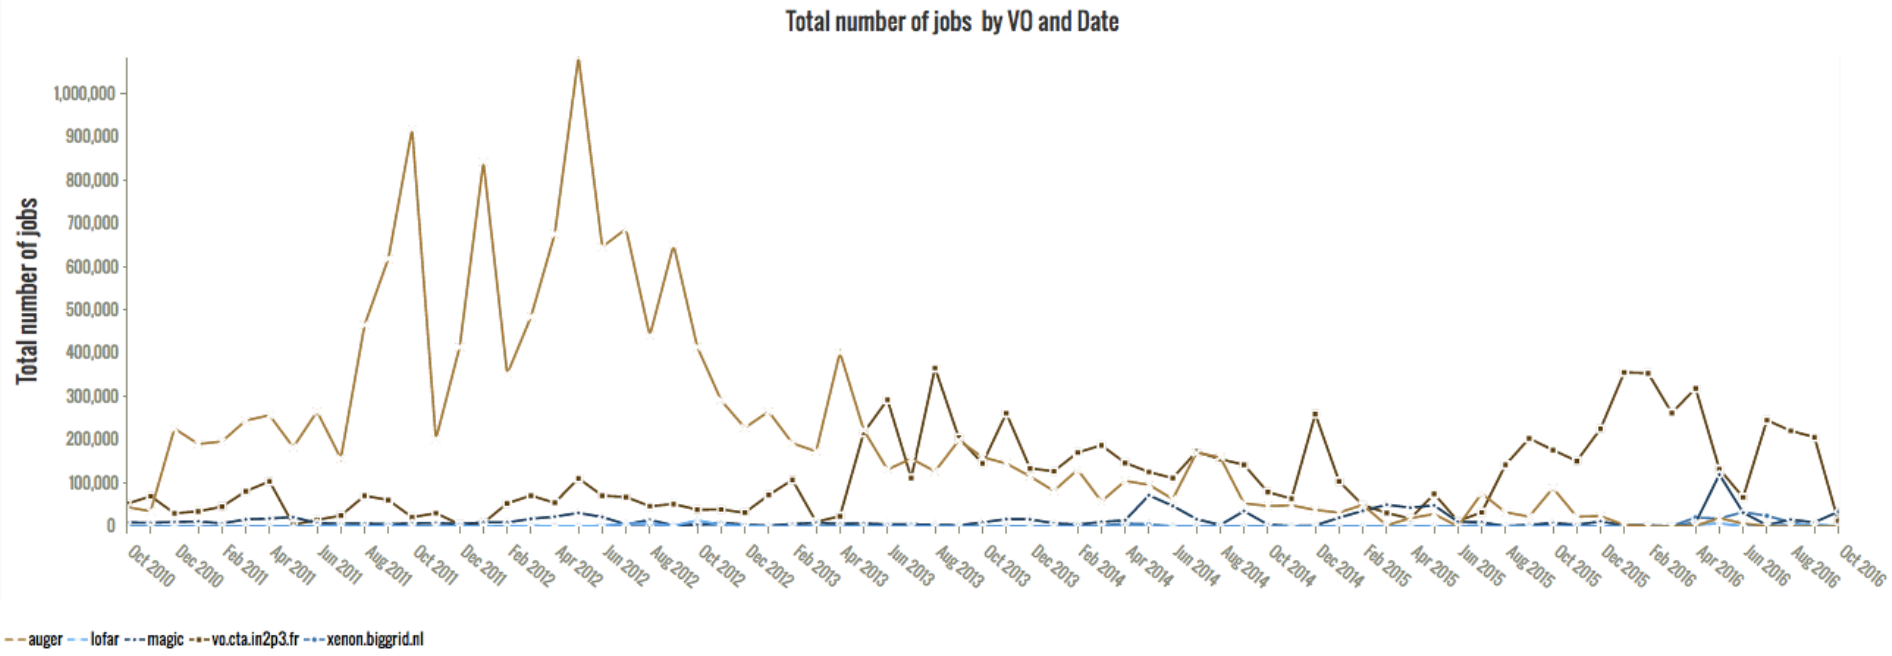

Supporting international research communities and thematic services

Installed compute capacity trends 2011-2016

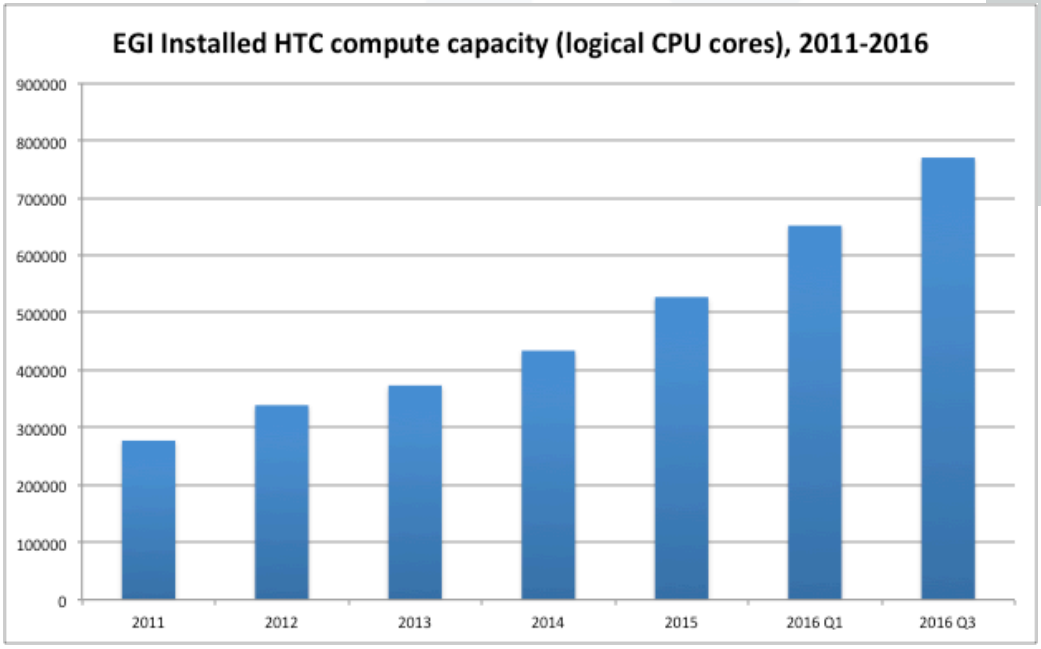

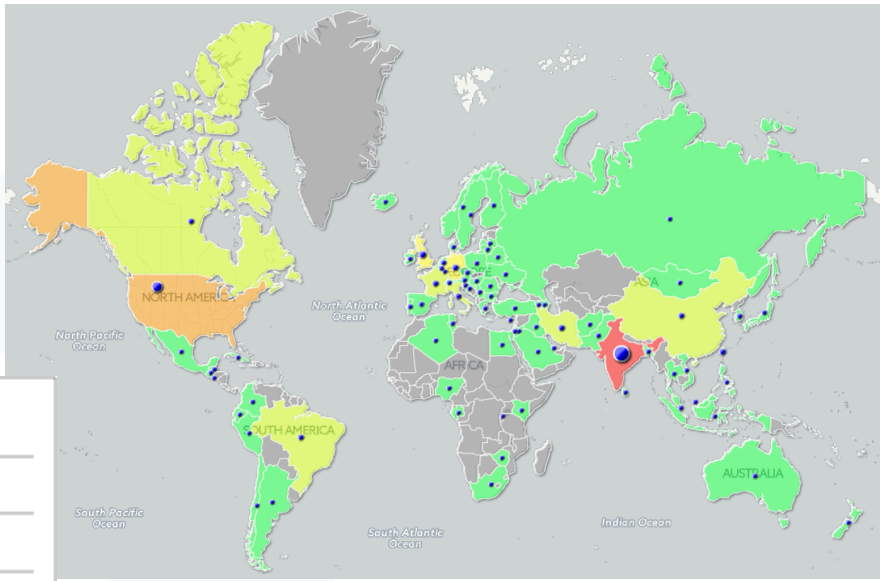

Structural Biology
Distribution of users (2016, QR3)

- 2700 users
- 81 countries

(credits: A. Bonvin, WeNMR)

Services Catalogue

Compute

Cloud Compute >

Run virtual machines on demand with complete control over computing resources

Cloud Container Compute >

Run Docker containers in a lightweight virtualised environment

High-Throughput Compute >

Execute thousands of computational tasks to analyse large datasets

Storage and Data

Online Storage >

Store, share and access your files and their metadata on a global scale

Archive Storage >

Back-up your data for the long term and future use in a secure environment

Data Transfer >

Transfer large sets of data from one place to another

Training

FitSM training >

Learn how to manage IT services with a pragmatic and lightweight standard

Training infrastructure >

Dedicated computing and storage for training and education

http://go.egi.eu/ServiceCatalogue

EGI-Engage

Cloud Compute

Run virtual machines on-demand with complete control over the computing resources

- On-demand provisioning
- Full control over computing resources
- Standard interface to deploy on multiple service providers

Benefits

- Execute compute- and data-intensive workloads, including GPGPU computing in the cloud
- Host long-running services
- Create disposable testing and development environments
- Select virtual machine configurations and application environments
- Manage your Cloud Compute resources

Communities using Cloud Compute

The EXTraS project is harvesting 13 years of data collected on-board

the ESA's X-ray space observatory XMM-Newton. The project is using Cloud Compute to implement four lines of analysis with ad-hoc software pipelines

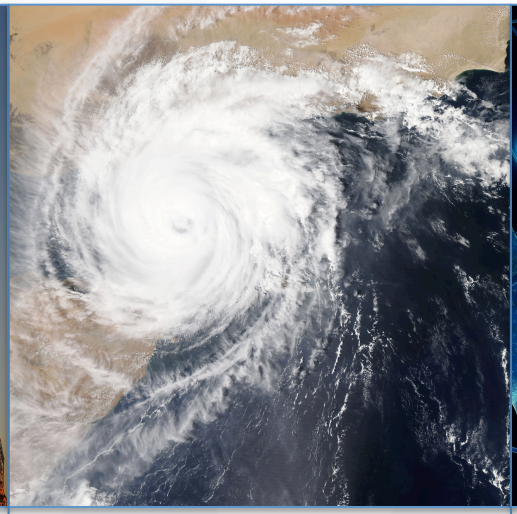

The DRIHM project is prototyping an e-infrastructure to simulate extreme hydro-meteorological events such as ash ooding.

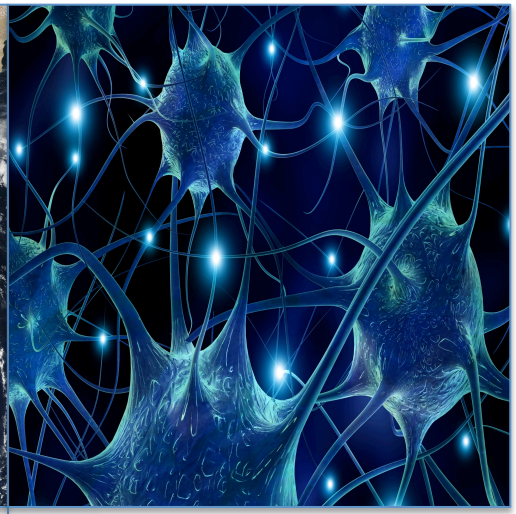

The National Bioinformatics Infrastructure of Sweden uses Cloud Compute to provide bioinformatics tools to their researchers, including tools to predict 3D protein structures, for example. So far, more than 6,700 unique users in 73 countries have made the most of these resources

EGI-Engage

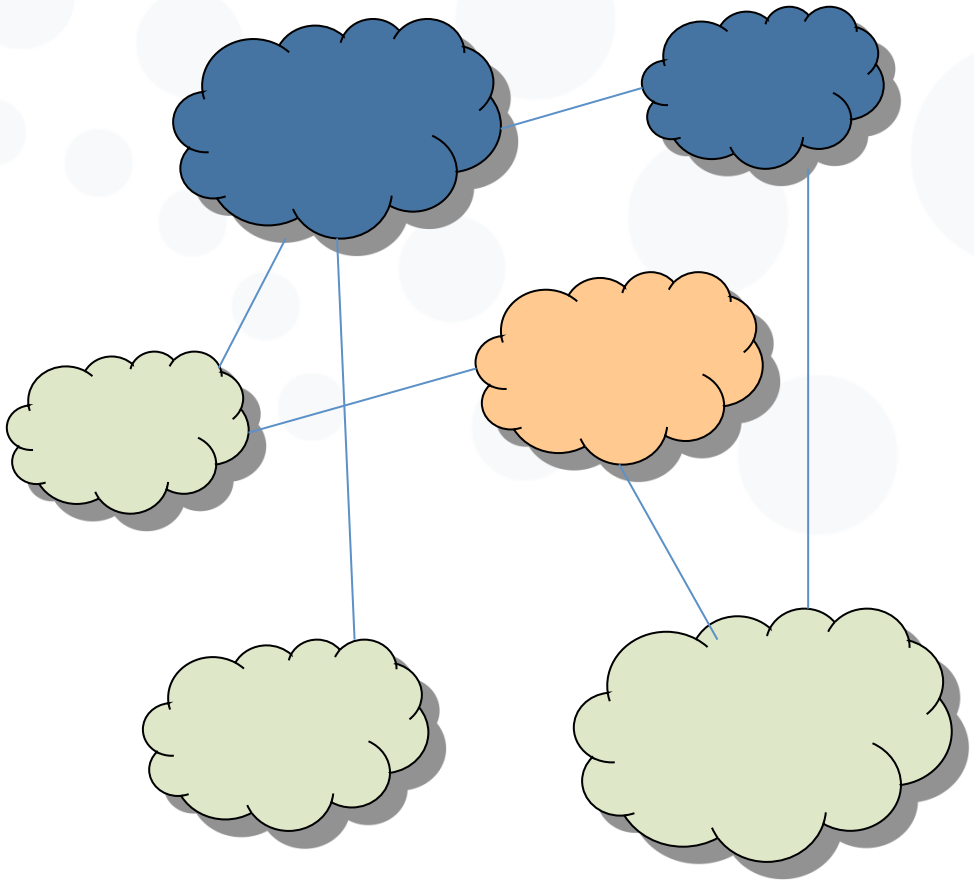

EGI Federated Cloud

- System of cloud infrastructures
- Standard user interfaces
  - Clouds and their interconnections are based on open standards, open technologies
  - Based on OCCI/OGF and OpenStack
- Harmonised operational behaviour
- Value proposition: distributed cloud computing for analysis of distributed large datasets

Benefits

VM and block storage management:

Occi - On every site

OpenStack Nova - On OS sites

Object storage management (optional):

CDMI - on any site

OpenStack SWIFT – on OS sites

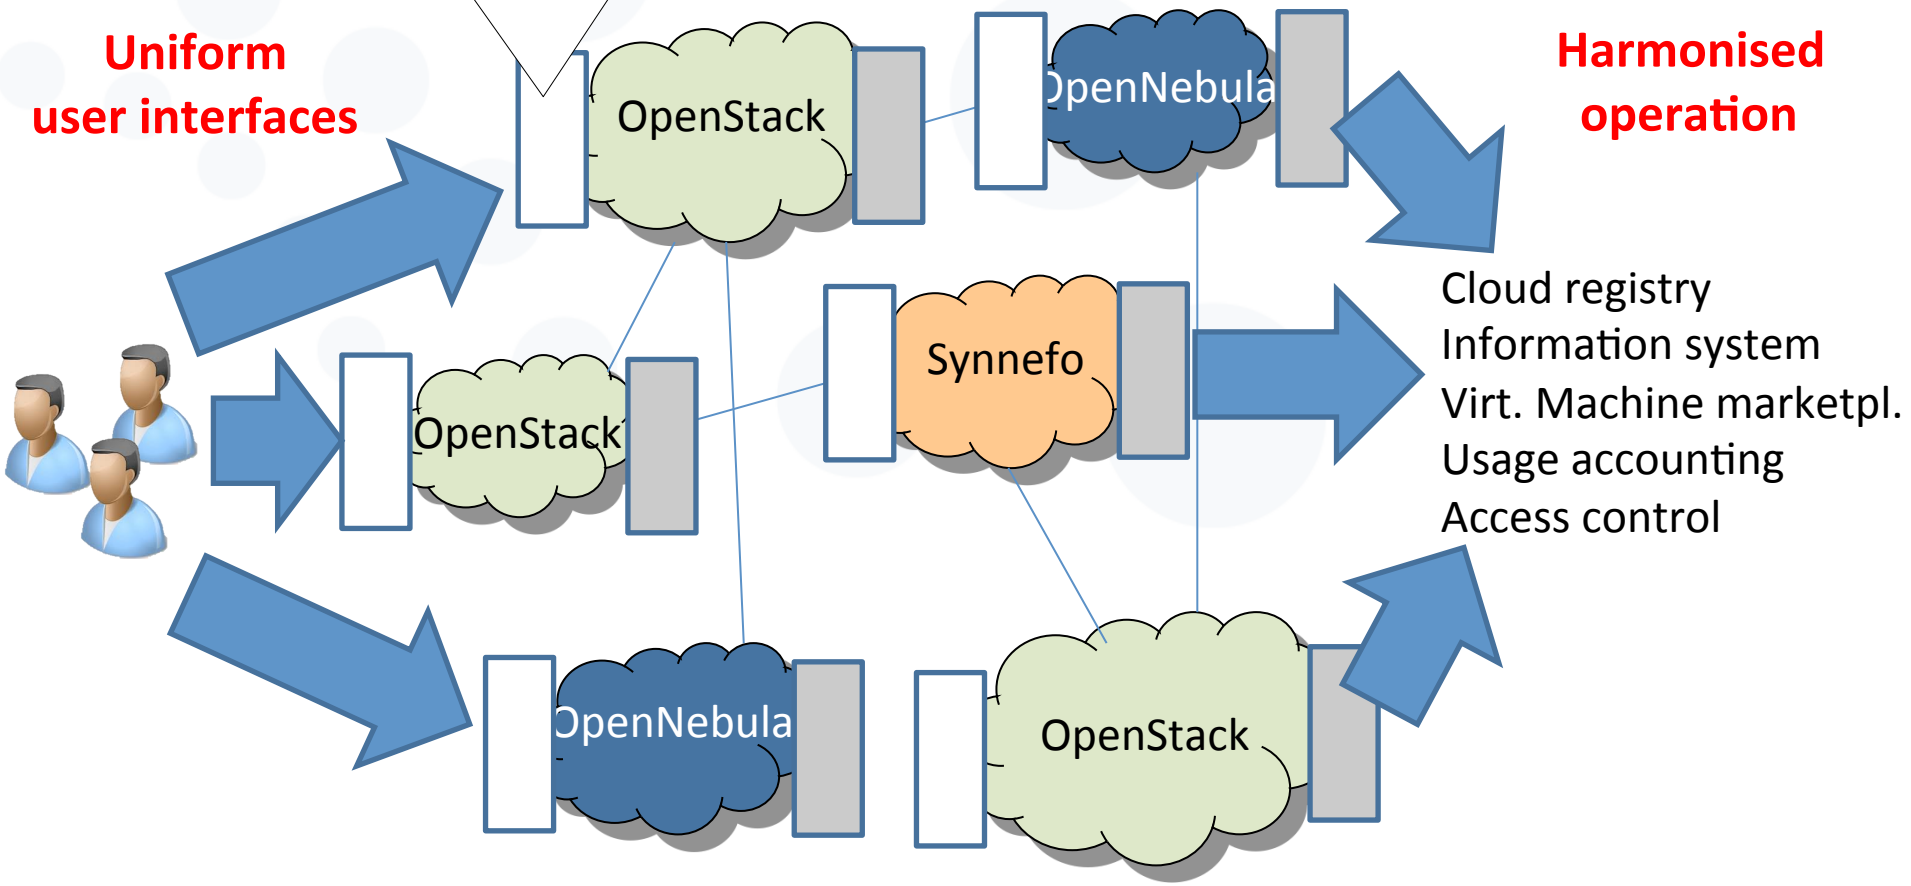

EGI Federated Cloud

EGI-Engage

EGI Federated Cloud is a collaboration of communities developing, innovating, operating and using cloud federations for research and education.

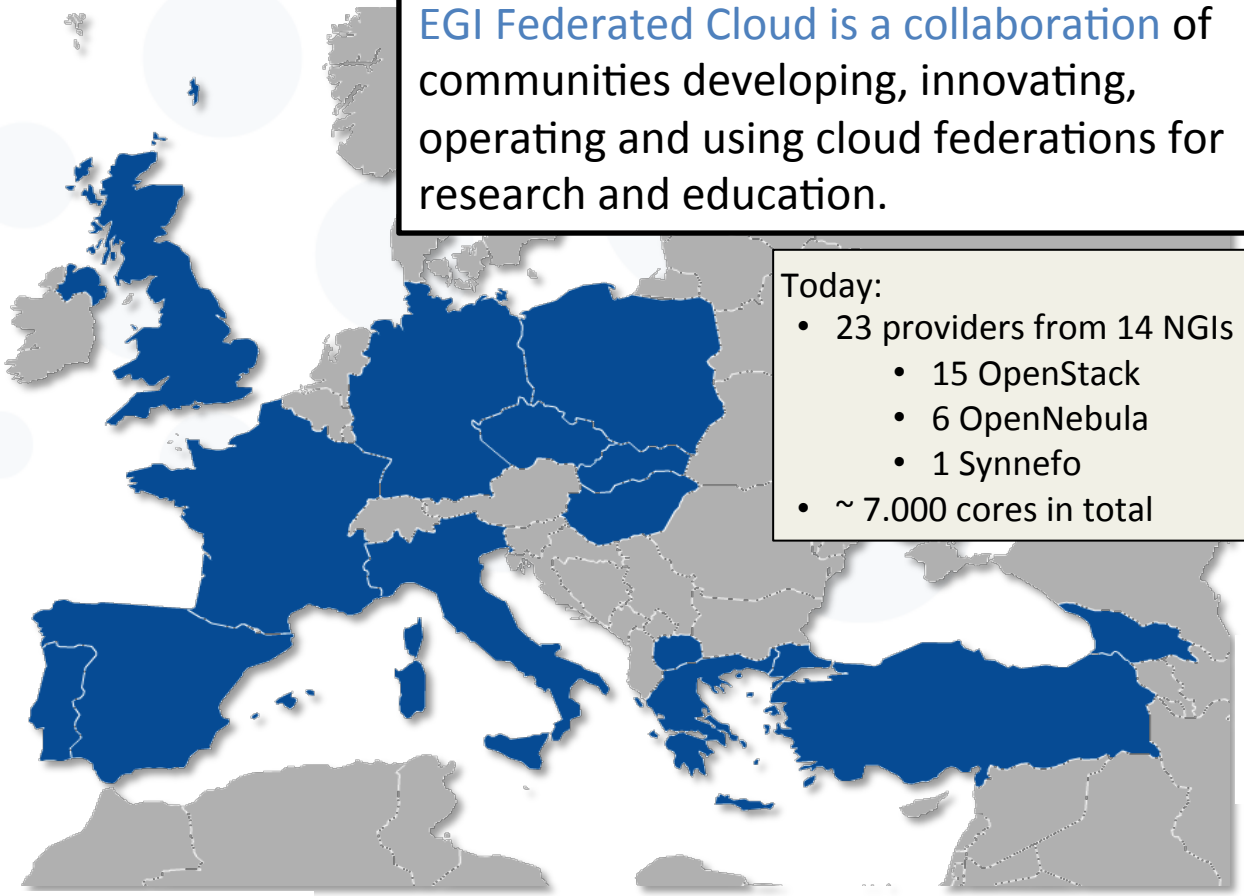

Today:

- 23 providers from 14 NGIs
  - 15 OpenStack
  - 6 OpenNebula
  - 1 Synnefo
- ~ 7.000 cores in total

Porting of LOFAR calibration pipeline 1/2

LOFAR use case details

- Measurement sets: **datacubes** (3D data): two Fourier spatial coordinate axes plus a spectral axis.
- A datacube can reach several **terabytes**.
- LOFAR telescope allows **up to 488 subbands**, which can reach several GBs.
- Each subband processed independently.

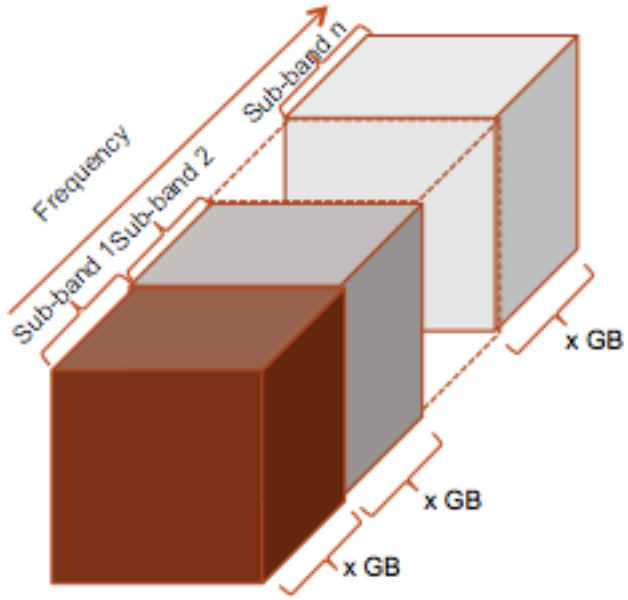

Credits: Susana Sánchez Expósito - CSIC

Porting of LOFAR calibration pipeline 2/2

Implementing the Use Case: COMPSs application

- COMPSs:
  - a data-driven programming model
    - it exploits the inherent parallelism of the applications
    - It executes the application tasks as soon as their input data are ready
  - a VM orchestrator
    - It starts and contextualize the VM instances needed to execute the application tasks
    - It also checks the status, gathers the outputs and deletes the VM instances
- Our COMPSs application:
  - A python script
  - It iterates over the subbands executing for each one a COMPSs task
  - They calls the LOFAR software (= executes a script) to process the subband.

```
import subprocess
import sys
import os
from pycompss.api.task import task
from pycompss.api.parameter import *

@task( script_name = FILE)
def iter_calib(script_name):
   os.chmod(script_name,0744)
   subprocess.call(script_name)
   print "end executiong"
if __name__ == "__main__":
   args = sys.argv[1:]
   DATA_PATH=args[0]
   TEMPLATE_FILE=args[1]
   f=open(TEMPLATE_FILE,'r')
   content=f.read()
   f.close()
   list_f=os.listdir(DATA_PATH)
   for directory in list_f: # Iterate over the data inputs
      if os.path.isdir(DATA_PATH+"/"+directory):
        new_content=content.replace("INPUTDATAPATH",directory)
         script_name="job"+directory+".sh"
         f=open(script_name,"w")
         f.write(new_content)
         f.close()
         iter_calib(script_name)
```

EGI-Engage

Outcomes

- The computing capabilities fulfil the requirements from the use case
  - The memory and cpu needs depends on the specific pipeline step, and the EGI federated cloud allows to configure virtual machines with different capabilities.
- A better storage solution is needed
  - The user data are too large to be stored in the VM images. They should be stored in volumes easily mountable from several VMs and synchronized across different cloud providers.
- COMPSs facilitates the porting and deployment of the application

EGI-Engage

Cloud Container Compute

Run Docker containers within isolated user-space with no overhead

- On-demand provisioning
- Lightweight environment for maximized performance
- Standard interface to deploy on multiple service providers

Benefits

- Reduce time to production by removing friction between development and operations environments
- Interoperable and transparent

EGI-Engage

High-Throughput Compute

Analyze large datasets by executing large numbers (thousands) of computational tasks

- Access to high-quality computing resources
- Integrated monitoring and accounting tools to provide information about the availability and resource consumption
- Workload and data management tools to manage all computational tasks

Benefits

- Large amounts of processing capacity over long periods of time
- Faster results for your research
- Shared resources among users, enabling collaborative research

CTA: Monte Carlo simulations and Analysis

- Use of EGI High Throughput Compute, Online Storage, Archive Storage
- Use of DIRAC for Data Cataloguing and Workload Management

CTA resource pool in EGI:

- 20 sites with approx. 8000 CPU cores
- Disk: 1.3 PB in 6 sites (Online Storage)
- Tape: 400 TB in 3 sites (Archive Storage)

DIRAC for CTA:

- CTA-specific extension of DIRAC
- File Catalogue ~21 million files
- Computing tasks
  - 8 million so far; 2.6 PB processed
  - Data transformation tasks vs. User tasks

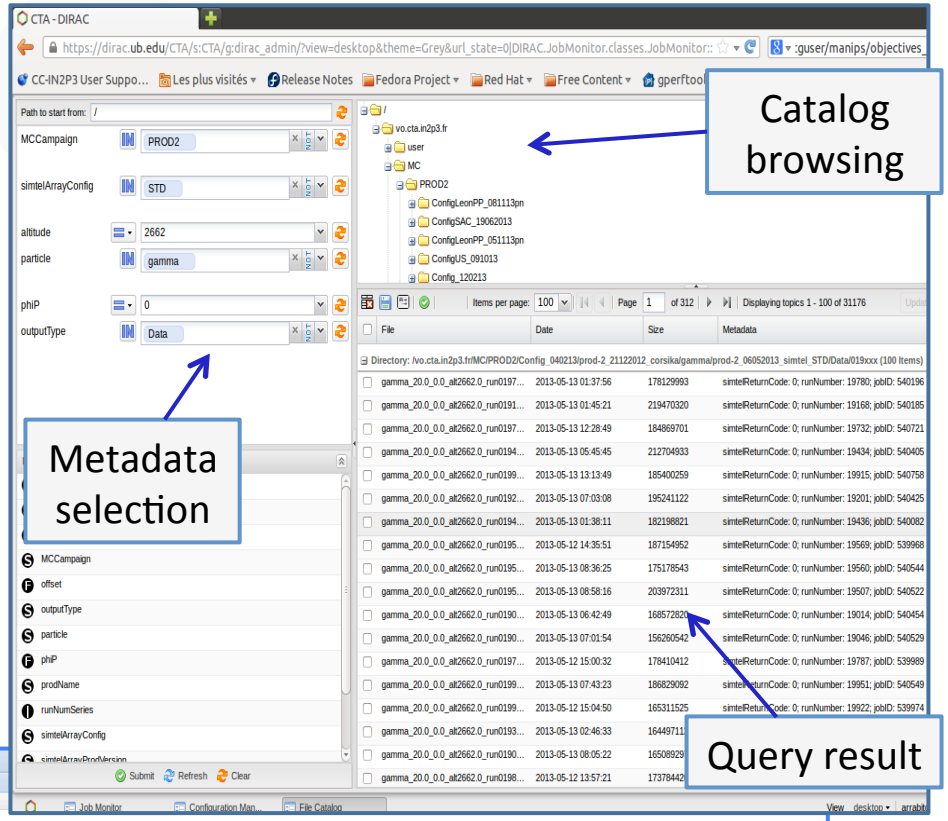

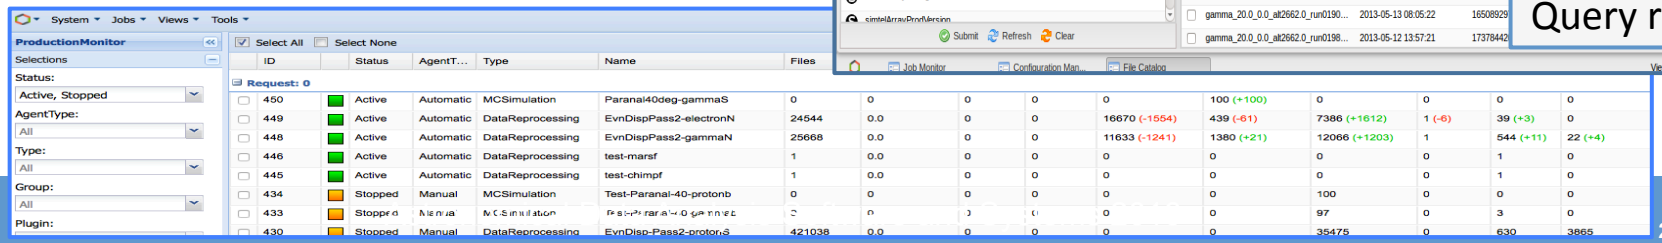

EGI-Engage

Online storage

Store, share and access your les and their metadata on a global scale

- Assign global identifiers to files
- Access highly-scalable storage from anywhere
- Control the data you share
- Organise your data using a exible hierarchical structure

Benefits

- Highly scalable storage system accessible from anywhere
- Easily share data
- Access through different interfaces

EGI-Engage

Archive storage

Back-up your data for the long term and future use in a secure environment

- Store data for long-term retention
- Store large amount of data
- Free up your online storage

Benefits

- Stores large amounts of data
- Long-term retention
- Reliable and interoperable

EGI-Engage

Data Transfer

Transfer large sets of data from one place to another

- Move research data fast
- Specialized analytics of on-going transfers
- User interface to manage transfer and network resources

Benefits

- Ideal for very large files
- Able to handle large amounts of files
- Transfer process with automatic retry

EGI AAI New Architecture

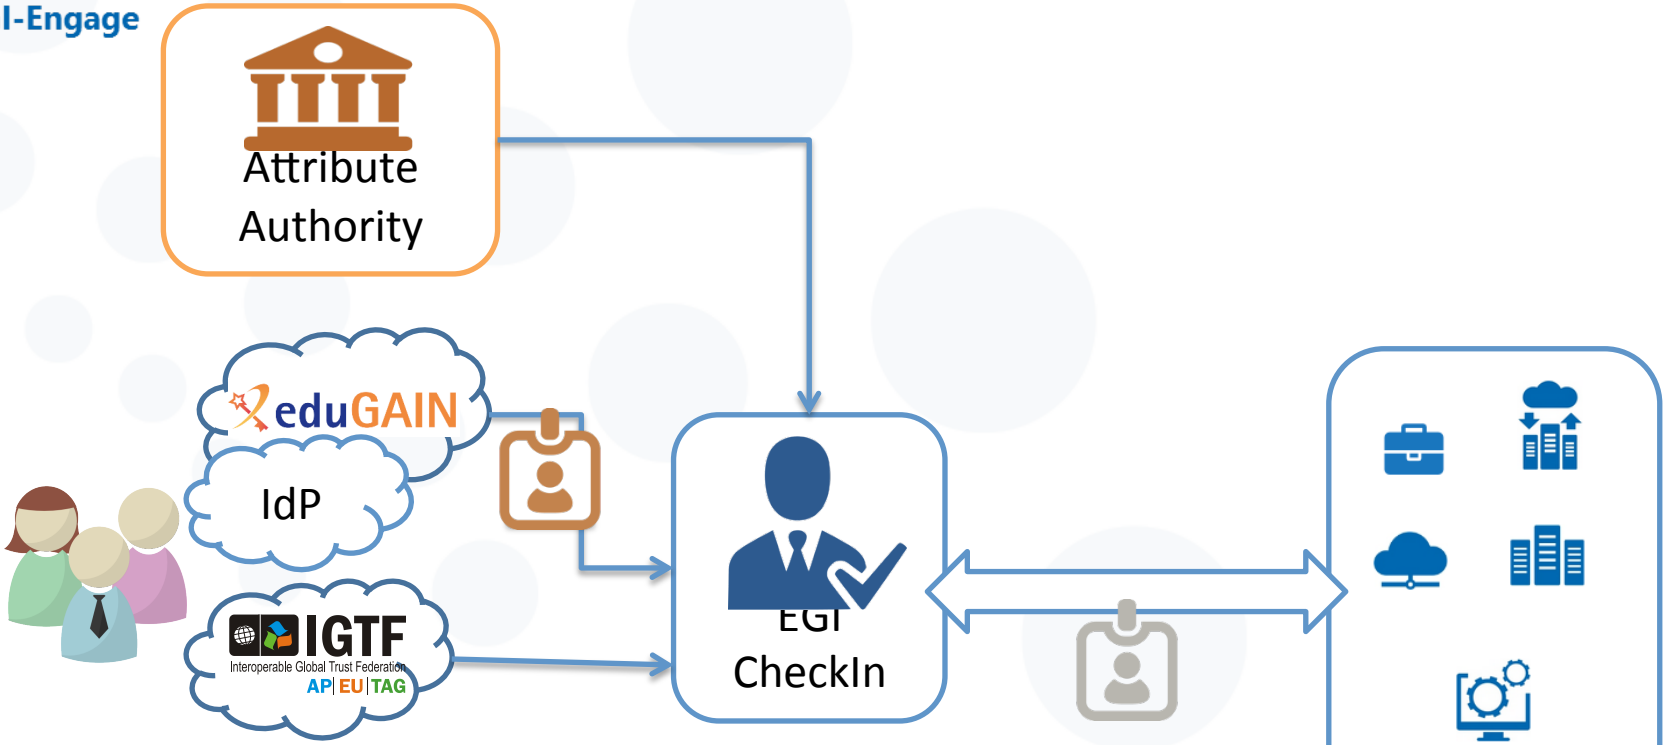

Mandatory Attributes		
	EGI UID	First name, last name
	email	affiliation

Why a IdP/SP Proxy?

- Service Providers (SPs) can have **one statically configured IdP**
- **No need to run an IdP Discovery Service** on each EGI SP
- Connected SPs get **harmonised user identifiers and accompanying attribute sets** from one or more AAs that can be interpreted in a uniform way for authZ purposes
- External IdPs only deal with a **single EGI SP** proxy

EGI services will not have to deal with the complexity of multiple IdPs/Federations/Attribute Authorities/ technologies.

EGI AAI Features

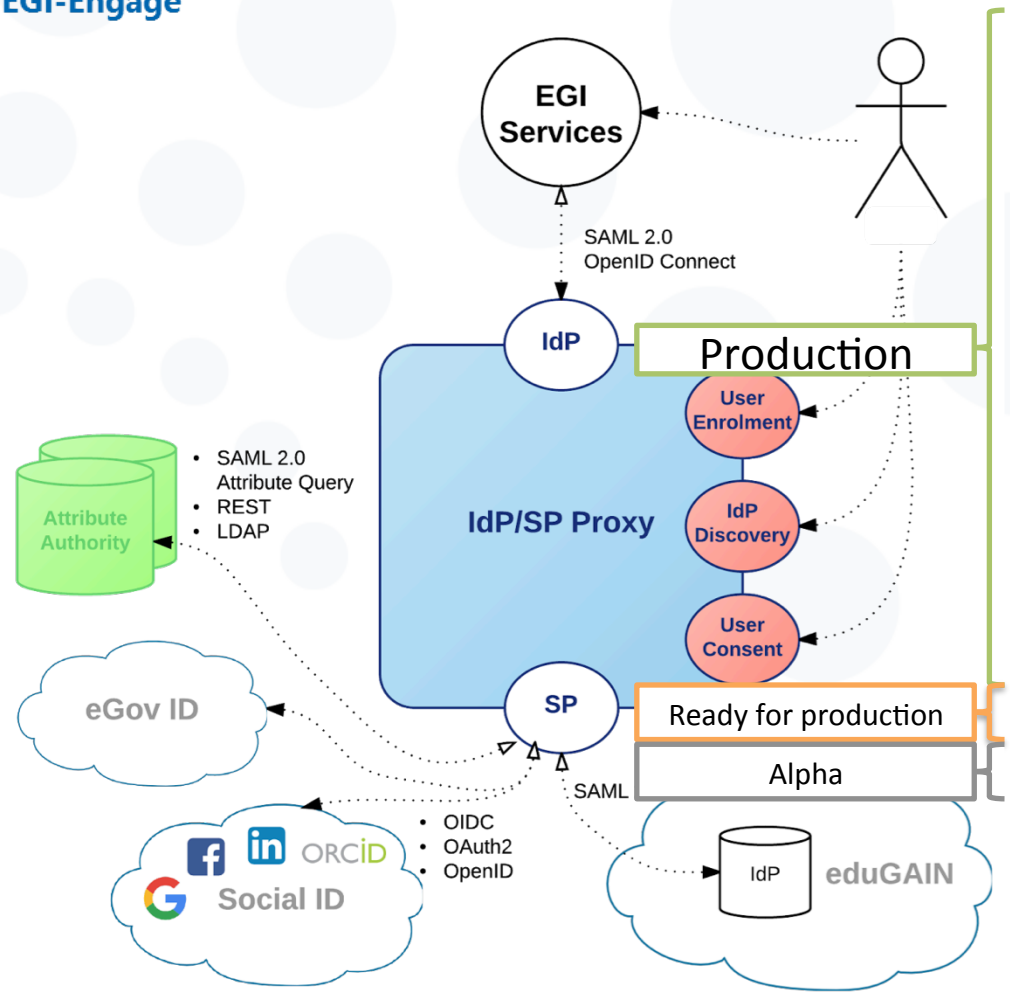

- Available via eduGAIN
- IdP Discovery
- User Enrolment
- User Consent
- Support for LoA
- Attribute Aggregation
  - SAML2.0 Attribute Query, REST, LDAP
- Support for OIDC/OAuth2 Providers
  - Google, Facebook, LinkedIn, ORCID
- Support for OIDC/OAuth2 services
- Experimental support for eIDAS

Open Data Platform

- Manage entire data life cycle from raw data to preservation
- Combine efficient computation services with open data managed by federated infrastructures
  - No local staging of data for processing
- Share public datasets for download or reuse
- Make public datasets discoverable

Open Data Platform interfaces

GUI

- Web based
- Easy data management and sharing, access control
- Publication of data items and collections

REST

- Advanced data and collection management API for integration with community tools and portals

CDMI

- Standard data management operations
- Advanced metadata queries
- Integration with future data management applications

POSIX

- Enable direct mounting of spaces in the local filesystem without full data transfer

OAI-PMH

- OAI Data Provider interface
- Dublin Core metadata by default
- More complex metadata can be registered in ODP manually

HTTP

- Direct download of open data from URL's

Open Data Platform workflow

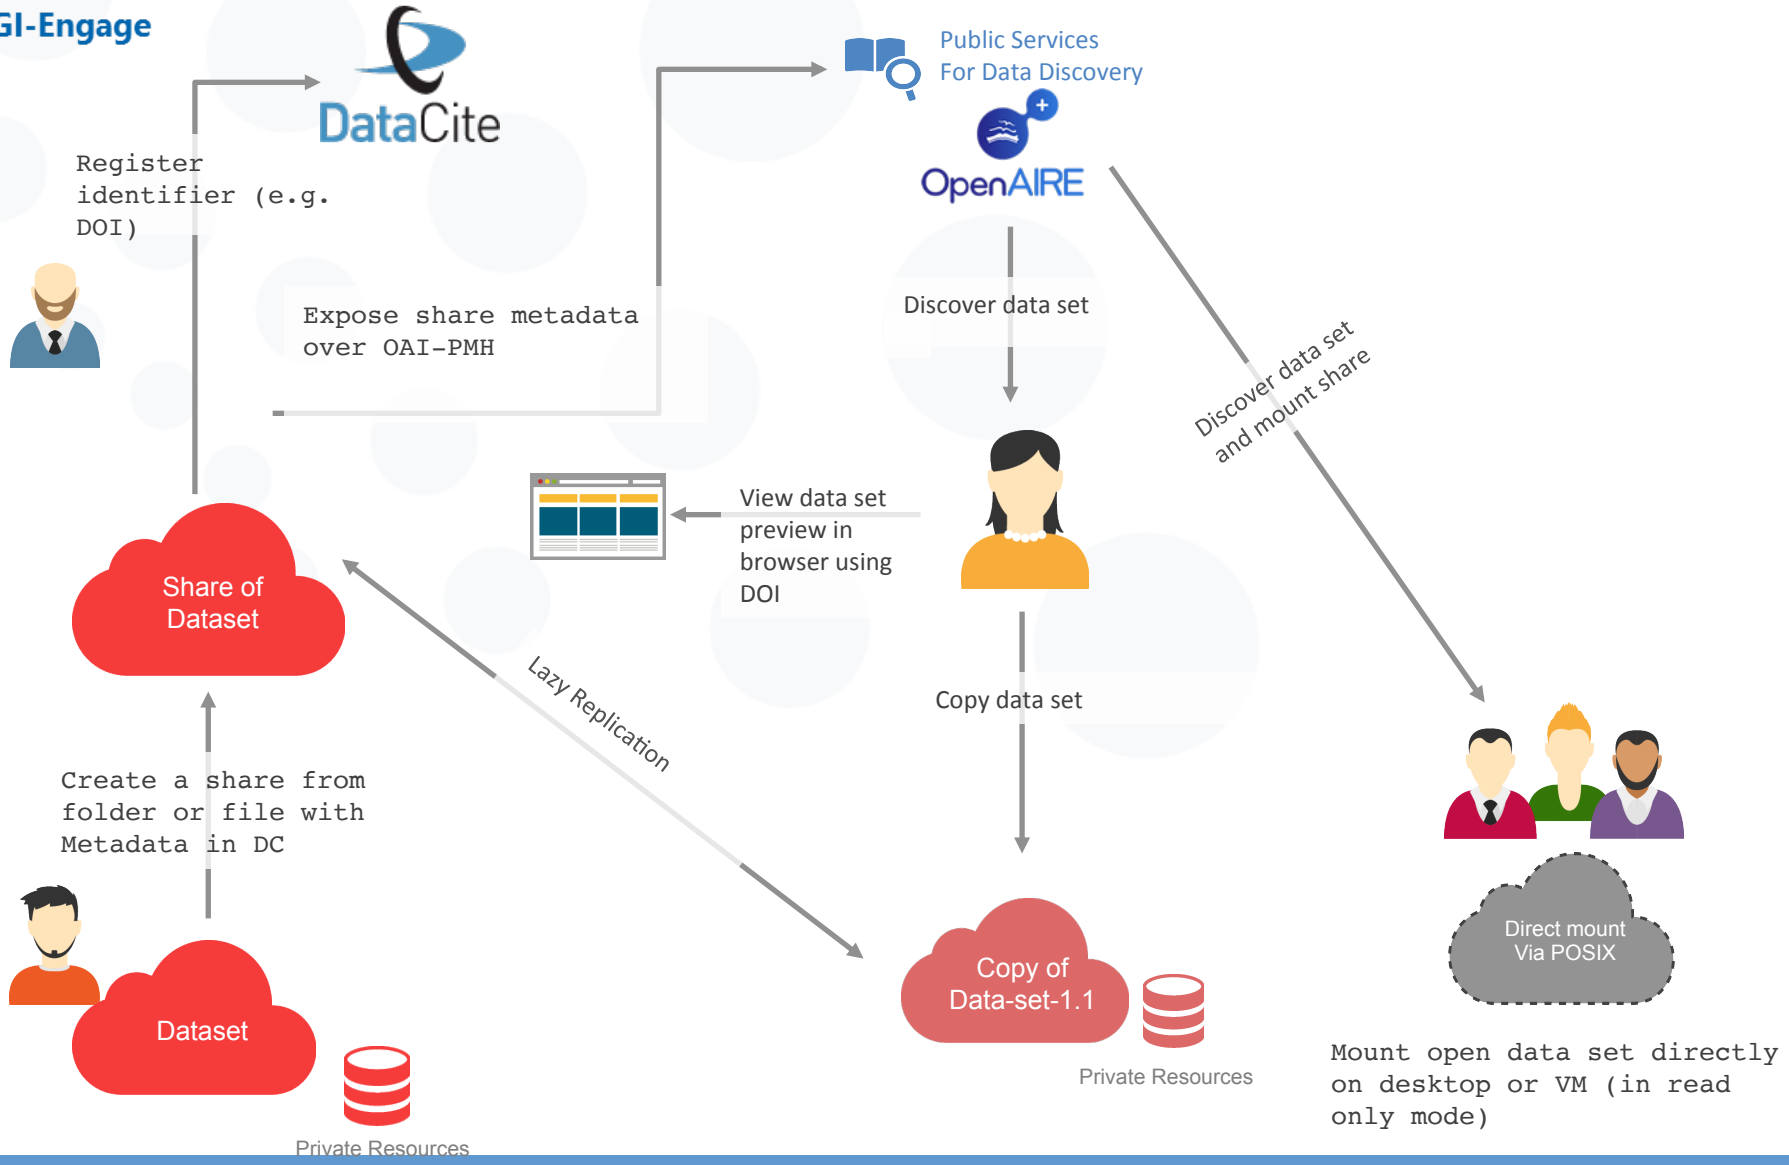

EGI role towards the European Open Science Cloud

- An Open Science Service Exchange as partnership of public/commercial organizations and initiatives responsible for
  - Provisioning of wide set of services to researcher and innovators → consolidation of national e-Infrastructures, open standards, technical and business process integration among the suppliers (e-Infrastructures, Research Infrastructures etc.)
  - Platform integration for community-specific capabilities with coordinated outreach
  - Aggregation of demand for economies of scale, technical requirements translations, cross-border access via brokering and procurement, end-to-end operations
  - Development of human capacity
  - A “Digital innovation hub” to support innovation with industry/SMEs

Thank you for your attention.

Questions?

Acknowledgements
This presentation used icons made by Freepik from www.flaticon.com