

OpenStack on the EGI Federated Cloud

Enol Fernández

Cloud Architect – EGI Foundation

enol.fernandez@egi.eu



EGI: Advanced Computing for Research

EGI's mission is to create and deliver open solutions for science and research infrastructures by federating digital capabilities, resources and expertise between communities and across national boundaries.

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



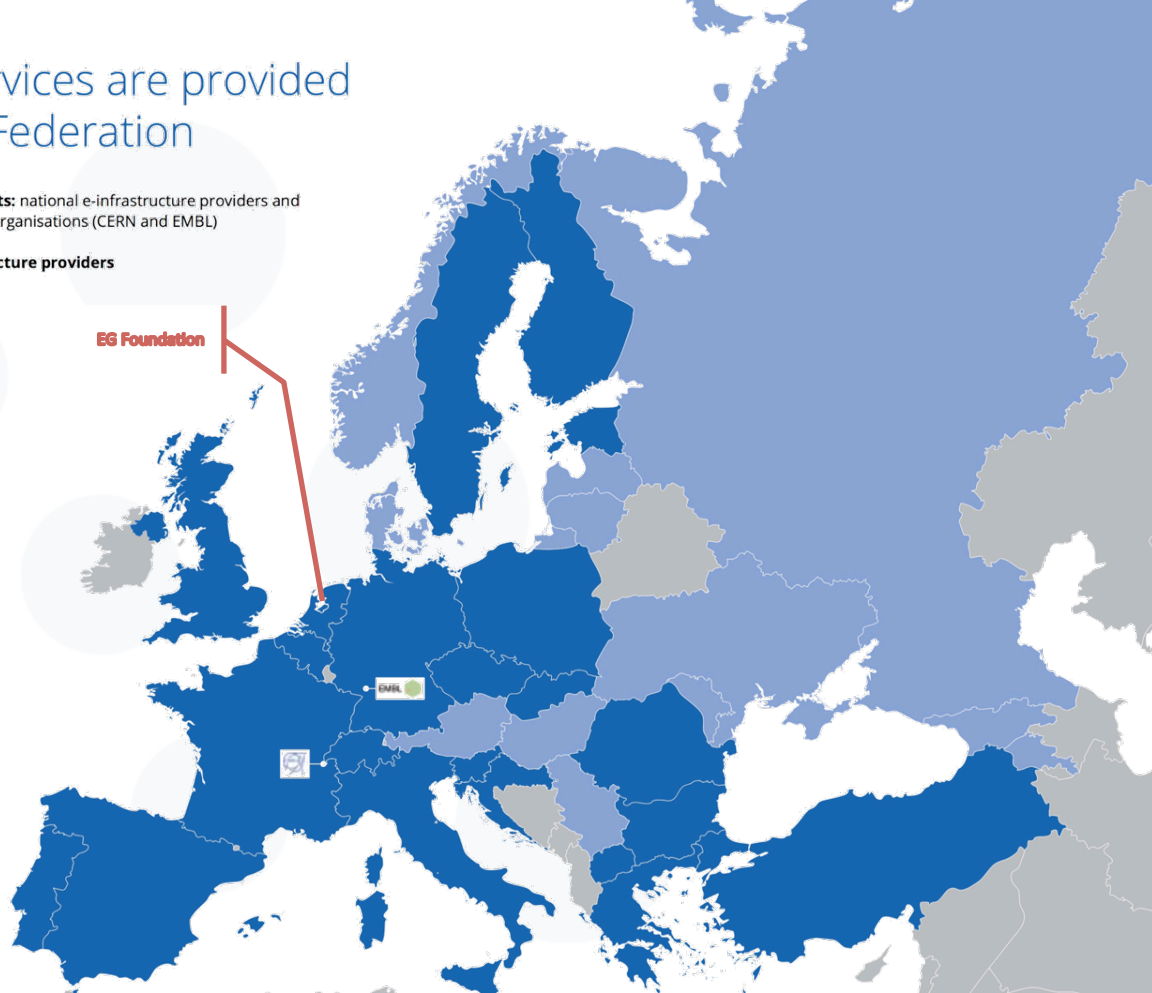
560,000
Terabytes of storage capacity



48,000
Users

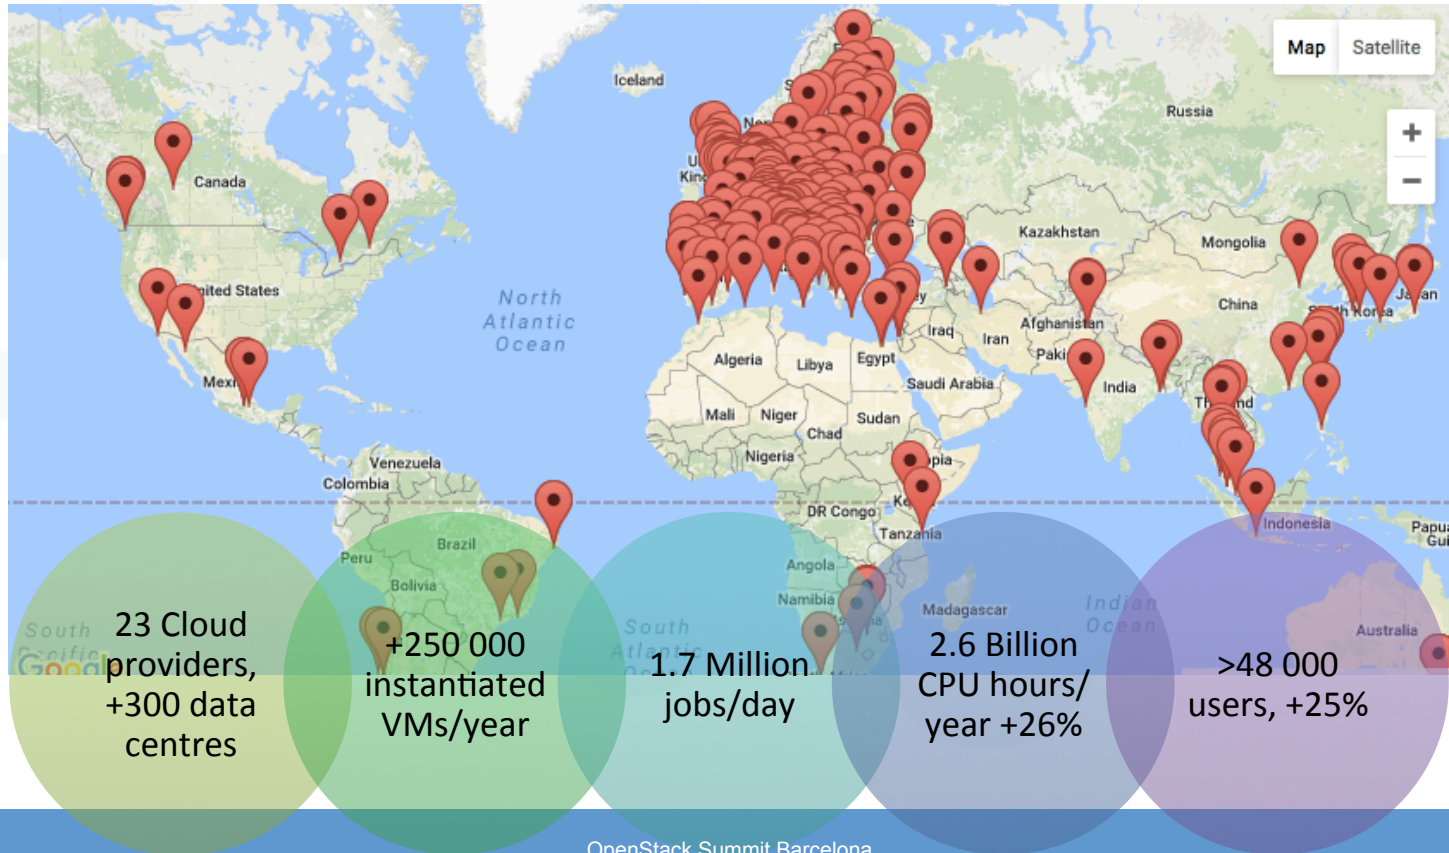


15
Research Infrastructures
Integrated with EGI



EGI Federation, 2016 QR3

The largest distributed compute e-Infra worldwide



International Partnerships



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

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

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



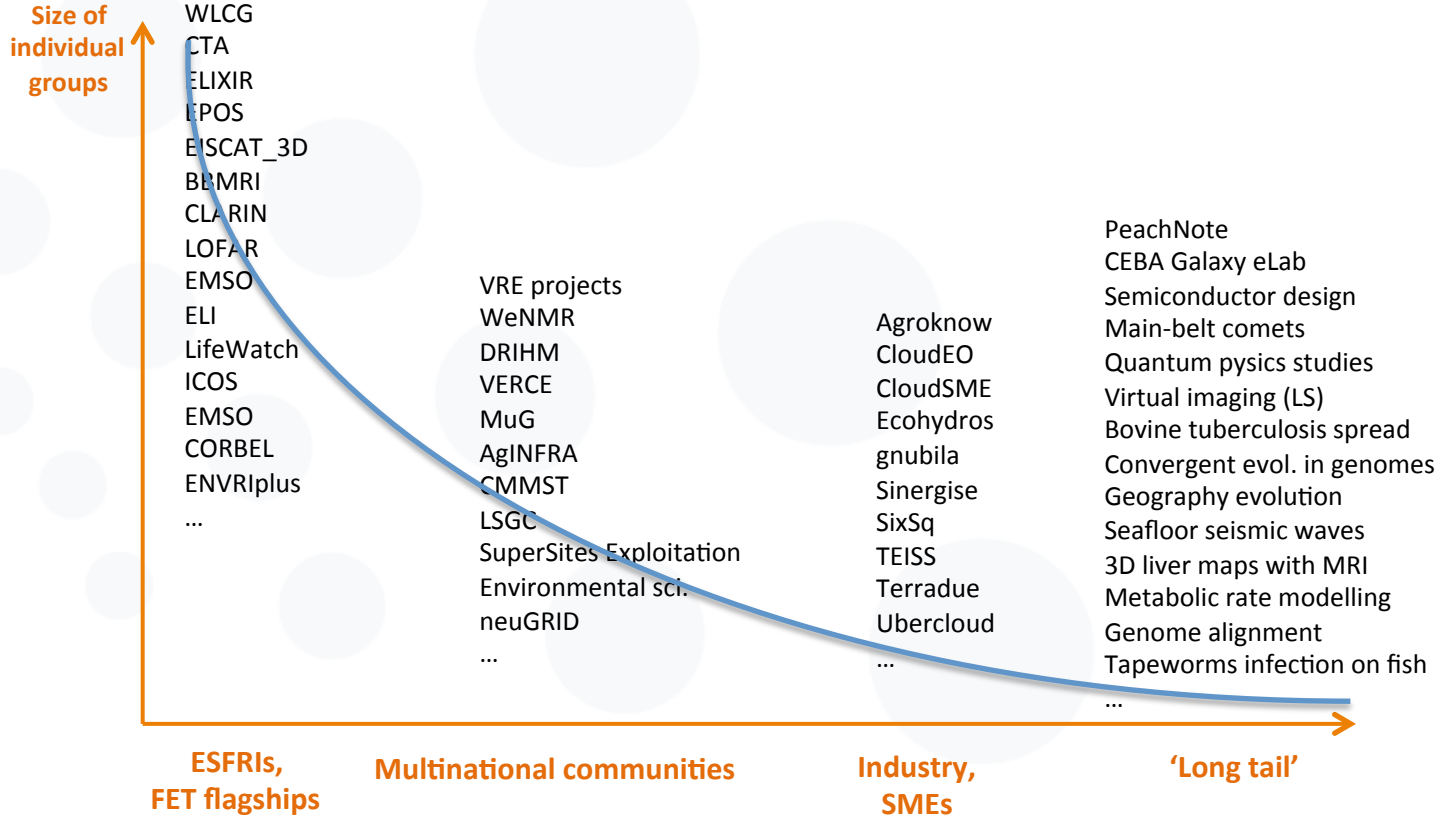
Data Transfer >

Transfer large sets of data from one place to another



<http://go.egi.eu/ServiceCatalogue>

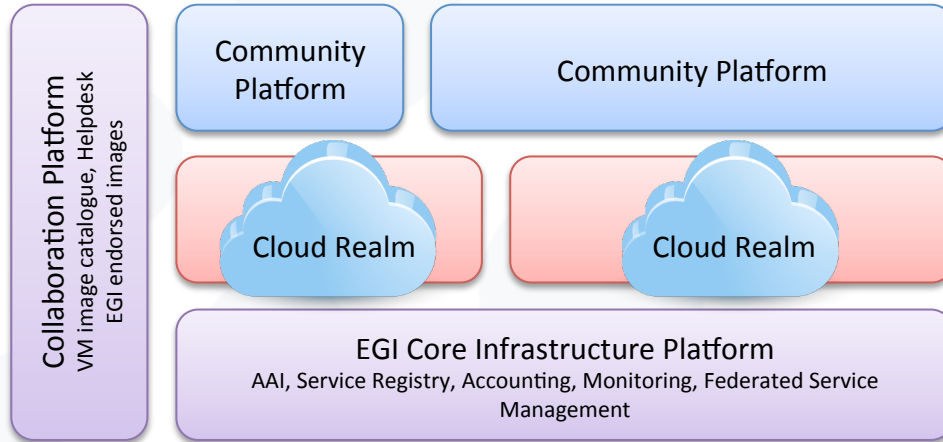
Serving researchers and innovators



EGI Federated Cloud

- EGI Federated Cloud is a collaboration of communities developing, innovating, operating and using cloud federations for research and education.
- 23 providers from 14 NGIs
 - 16 OpenStack
 - 6 OpenNebula
 - 1 Synnefo
- ~7K CPU cores





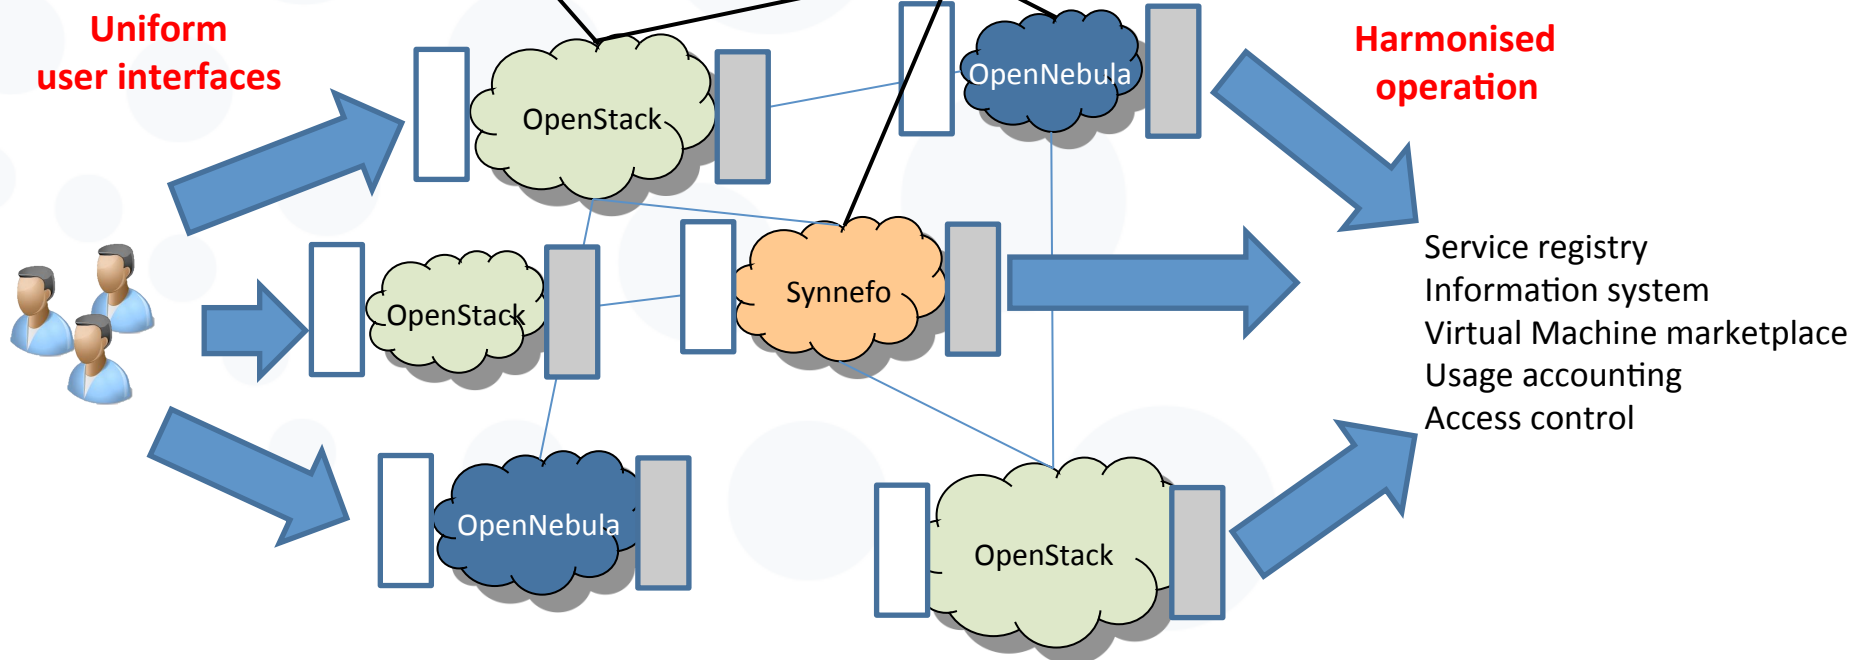
Cloud Realm

subset of cloud providers exposing homogeneous cloud management interfaces and capabilities which use the services of the EGI Core Infrastructure Platform for creating a federation

Community Platforms

provide community-specific data, tools and applications and can be supported by one or more realms.

Cloud Federation

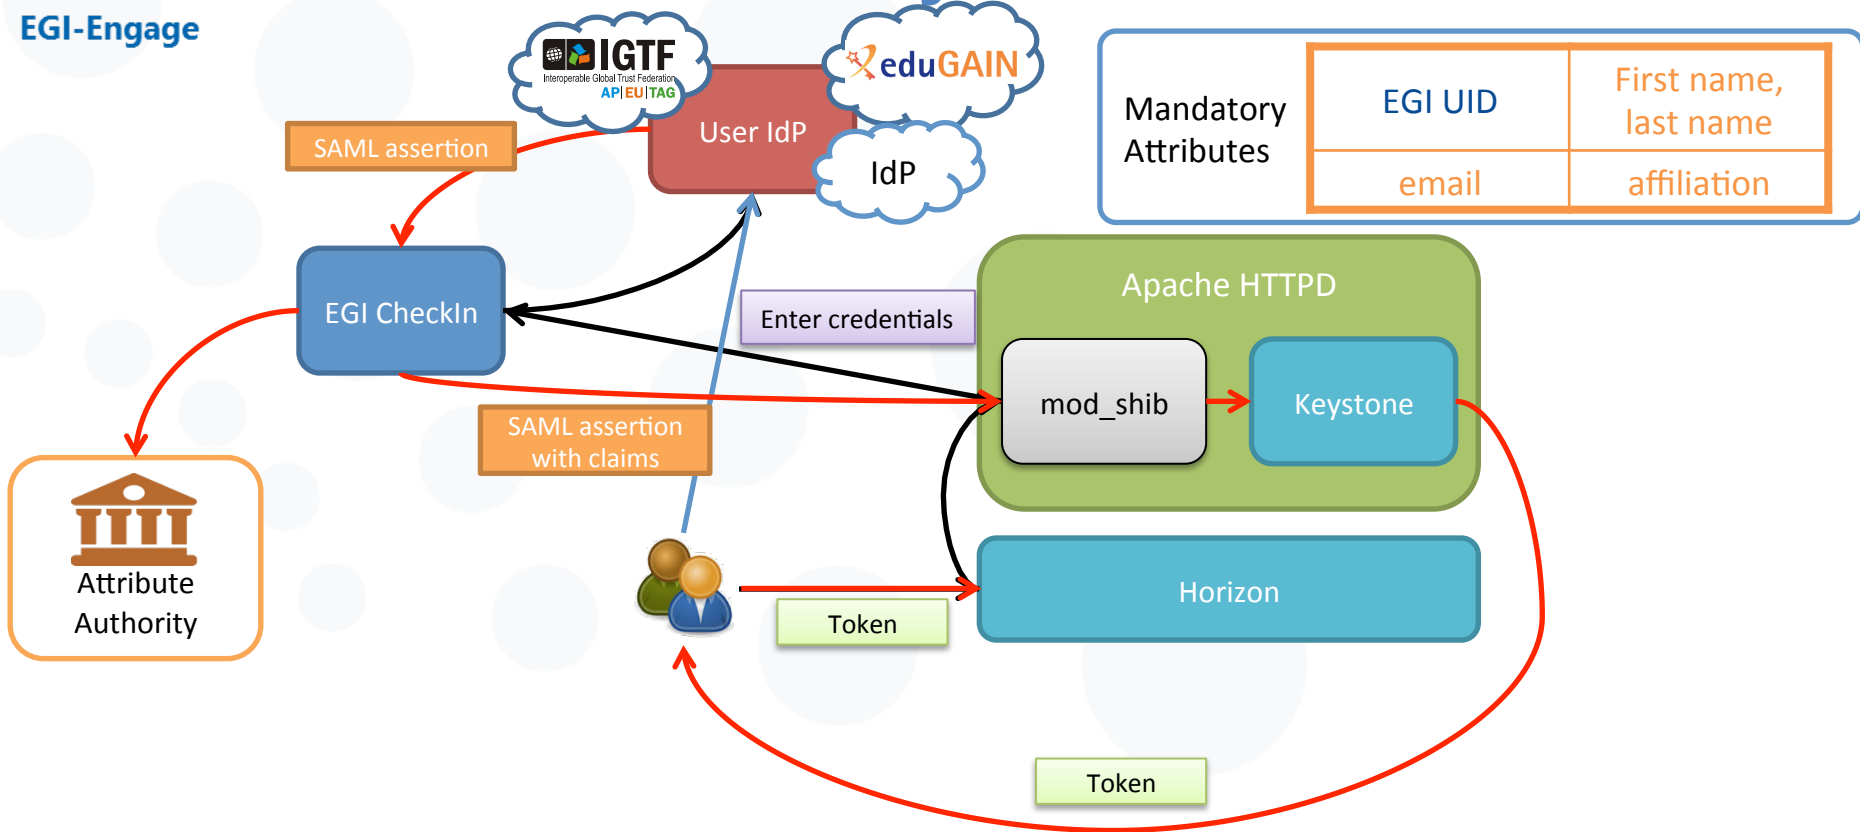


- Users identified with X.509 certificates (IGTF Federation) with VOMS extensions
 - VOMS (Virtual Organization Management System) provides attributes on membership to VOs, groups and roles on the VO
 - Not user-friendly, problematic in web-based GUIs
- Now in transition to new EGI AAI (**EGI CheckIn**)
 - Federated identity standards (SAML, OpenID Connect)
 - Allows users to authenticate with their institutional accounts
 - Integration with Attribute Authorities beyond VOMS

- WSGI filter for Keystone V2 API
- Extracts information from VOMS proxies to perform AuthN/AuthZ
- Can manage federation users
 - Add users to Keystone
 - Add roles to users in tenants
- Mapping VOMS → Keystone defined on file

<https://github.com/IFCA/Keystone-VOMS>

EGI AAI and Keystone OS-FEDERATION



Service Registry

- All Resource Centers must register their services at the EGI central catalog: <https://goc.egi.eu>
- Static information about services endpoints
 - org.openstack.nova and org.openstack.swift service types
- Web frontend and API access

GOCDB 5.6

Services
All Services in GOCDB

Filter (clear)

Service Type: org.openstack.nova NGI: (all)

Search for text in Hostname or Service Description:

Production Service: Y Monitored Service: (all) Site Certification: Certified

Service Scopes: Service Scopes Scope match: all (selected tags are AND'd)

Service Extension Name: (none) Extension Value:

Include Closed Sites: Filter Services


8 Services (Showing 1 - 50)

Hostname	Service Type	Production	Monitored	Host Site	Scope(s)
cloud.recas.ba.infn.it	org.openstack.nova	✓	✓	RECAS-BARI	EGI, lhcb, wlcg
controller.ceta-ciemat.es	org.openstack.nova	✓	✓	CETA-GRID	EGI, alice, wlcg
fsd-cloud.zam.kfa-juelich.de	org.openstack.nova	✓	✓	FZJ	EGI
sbgcloud.in2p3.fr	org.openstack.nova	✓	✓	IN2P3-IRES	EGI, alice, cms, tier2, wlcg
server4-epsh.unizar.es	org.openstack.nova	✓	✓	BIFI	EGI

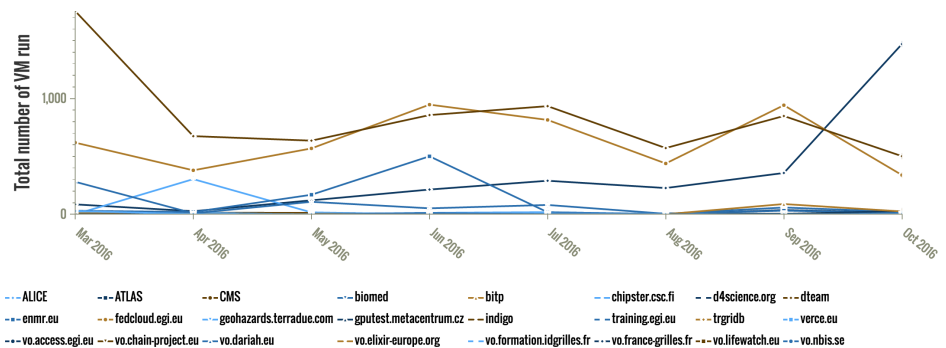
- Real-time information provided by BDII
 - Hierarchical information discovery system based on LDAP
 - Using standard [Glue Schema 2](#)
- Resource Centers publish actual capabilities
 - Available images & flavors
 - Supported user groups (VOs)
 - Available resources
- Cloud-bdii-provider
 - Gathers information from OS services using public APIs and puts it into Glue Schema
 - <https://github.com/EGI-FCTF/cloud-bdii-provider>

- Collect, aggregate and display **usage information** across the whole federation.
- **OGF Usage Record** extended for Cloud
- cASO produces accounting records using nova (and optionally ceilometer) APIs
 - <https://github.com/IFCA/caso>

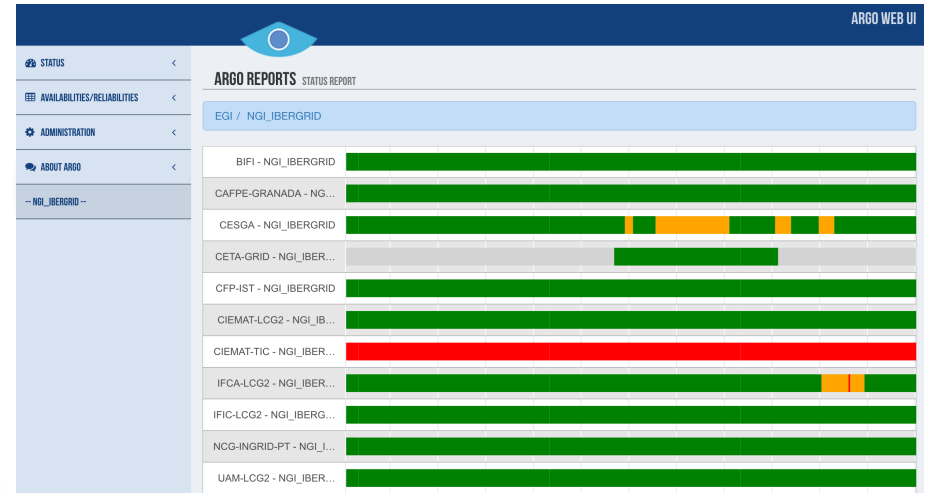
Total number of VM run by Region and Date



Total number of VM run by VO and Date

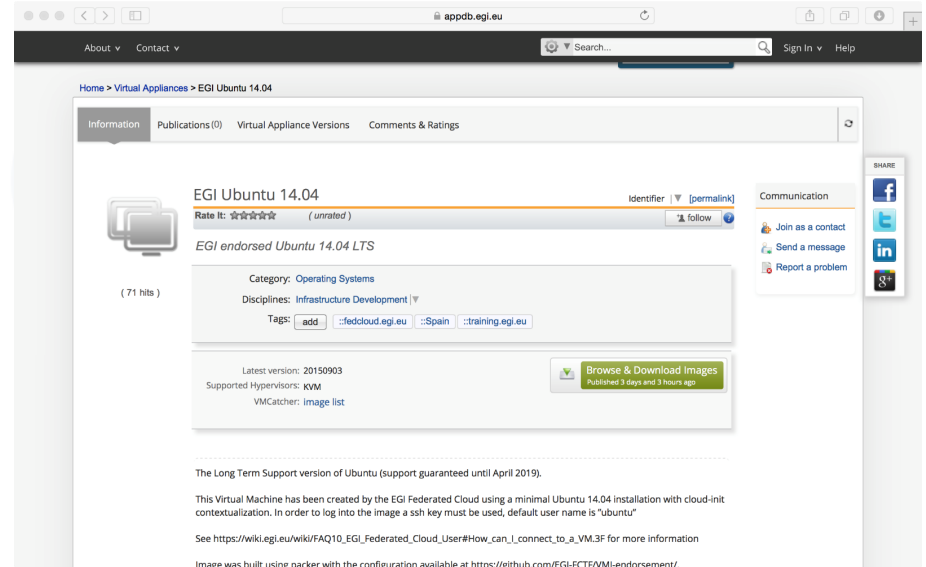


- Health monitoring of services
- Automatic discovery of services using GOCDB
- A/R metrics for SLA/OLAs
- Powered by EGI ARGO
 - <https://argoeu.github.io/>

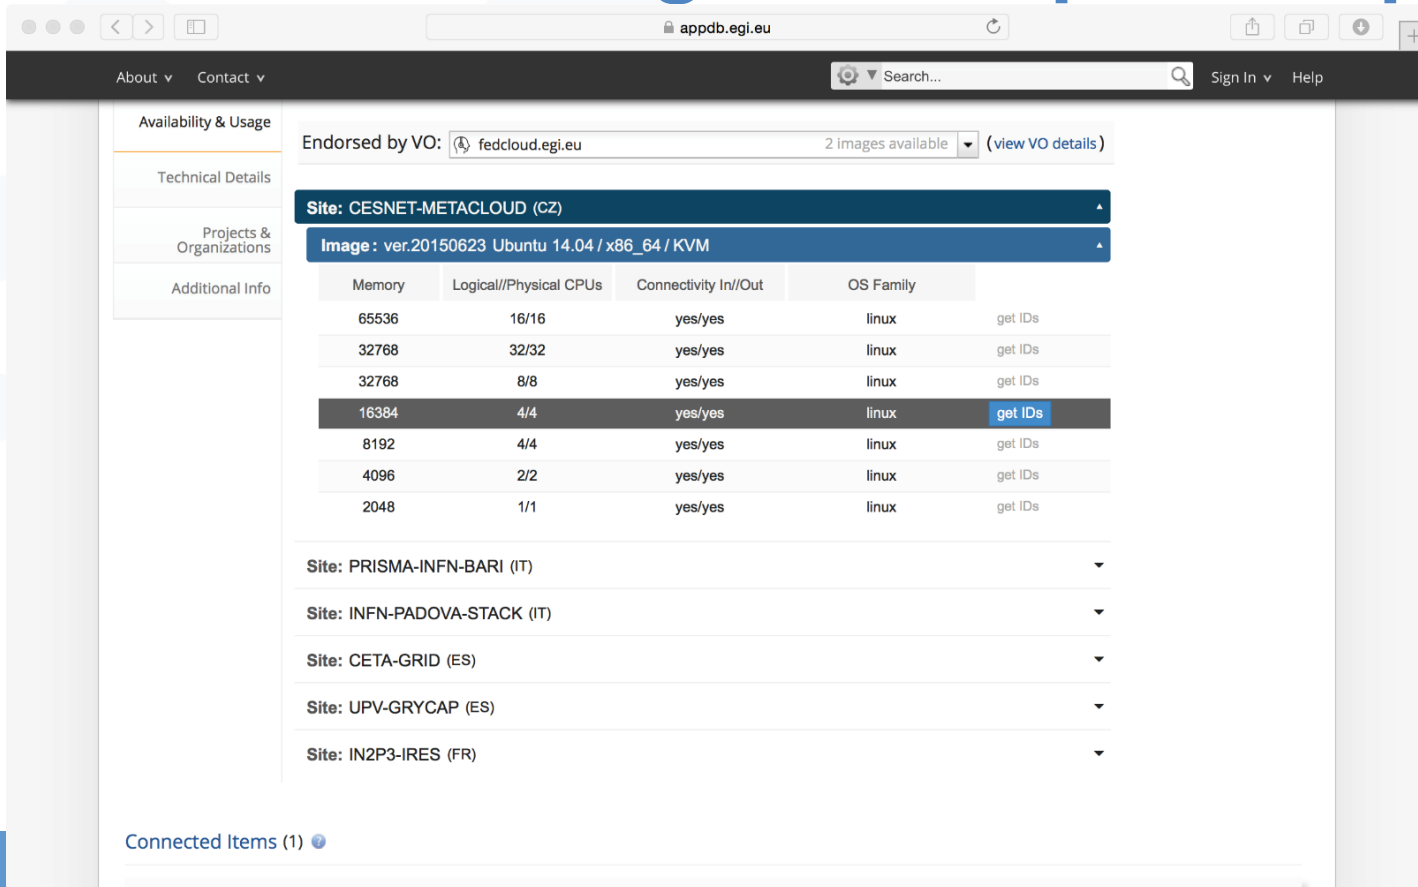


VM Image Marketplace: AppDB

- Open Library of Virtual Appliances
 - Use on clouds or for personal download
 - Re-use, share, associate contextualization
- EGI endorsed VM images, securely configured and tested
- Community curated sets of images
 - Automatic distribution of sets to cloud providers
 - <https://github.com/alvarolopez/atrope>



VM Image Marketplace: AppDB



The screenshot shows a web browser window at appdb.egi.eu. The page features a navigation bar with 'About' and 'Contact' menus, a search bar, and 'Sign In' and 'Help' links. A sidebar on the left contains menu items: 'Availability & Usage', 'Technical Details', 'Projects & Organizations', and 'Additional Info'. The main content area displays search results for 'fedcloud.egi.eu', showing '2 images available'. The first result is expanded to show details for 'Site: CESNET-METACLOUD (CZ)' and 'Image: ver.20150623 Ubuntu 14.04 / x86_64 / KVM'. Below this, a table lists various VM configurations with columns for Memory, Logical/Physical CPUs, Connectivity In/Out, OS Family, and a 'get IDs' link. The table has 7 rows, with the 4th row (16384 memory) highlighted. Below the table, several other sites are listed with dropdown arrows: PRISMA-INFN-BARI (IT), INFN-PADOVA-STACK (IT), CETA-GRID (ES), UPV-GRYCAP (ES), and IN2P3-IRES (FR). At the bottom left, there is a 'Connected Items (1)' link.

Endorsed by VO: 2 images available [\(view VO details\)](#)

Site: CESNET-METACLOUD (CZ)

Image: ver.20150623 Ubuntu 14.04 / x86_64 / KVM

Memory	Logical/Physical CPUs	Connectivity In/Out	OS Family	
65536	16/16	yes/yes	linux	get IDs
32768	32/32	yes/yes	linux	get IDs
32768	8/8	yes/yes	linux	get IDs
16384	4/4	yes/yes	linux	get IDs
8192	4/4	yes/yes	linux	get IDs
4096	2/2	yes/yes	linux	get IDs
2048	1/1	yes/yes	linux	get IDs

Site: PRISMA-INFN-BARI (IT) [▼](#)

Site: INFN-PADOVA-STACK (IT) [▼](#)

Site: CETA-GRID (ES) [▼](#)

Site: UPV-GRYCAP (ES) [▼](#)

Site: IN2P3-IRES (FR) [▼](#)

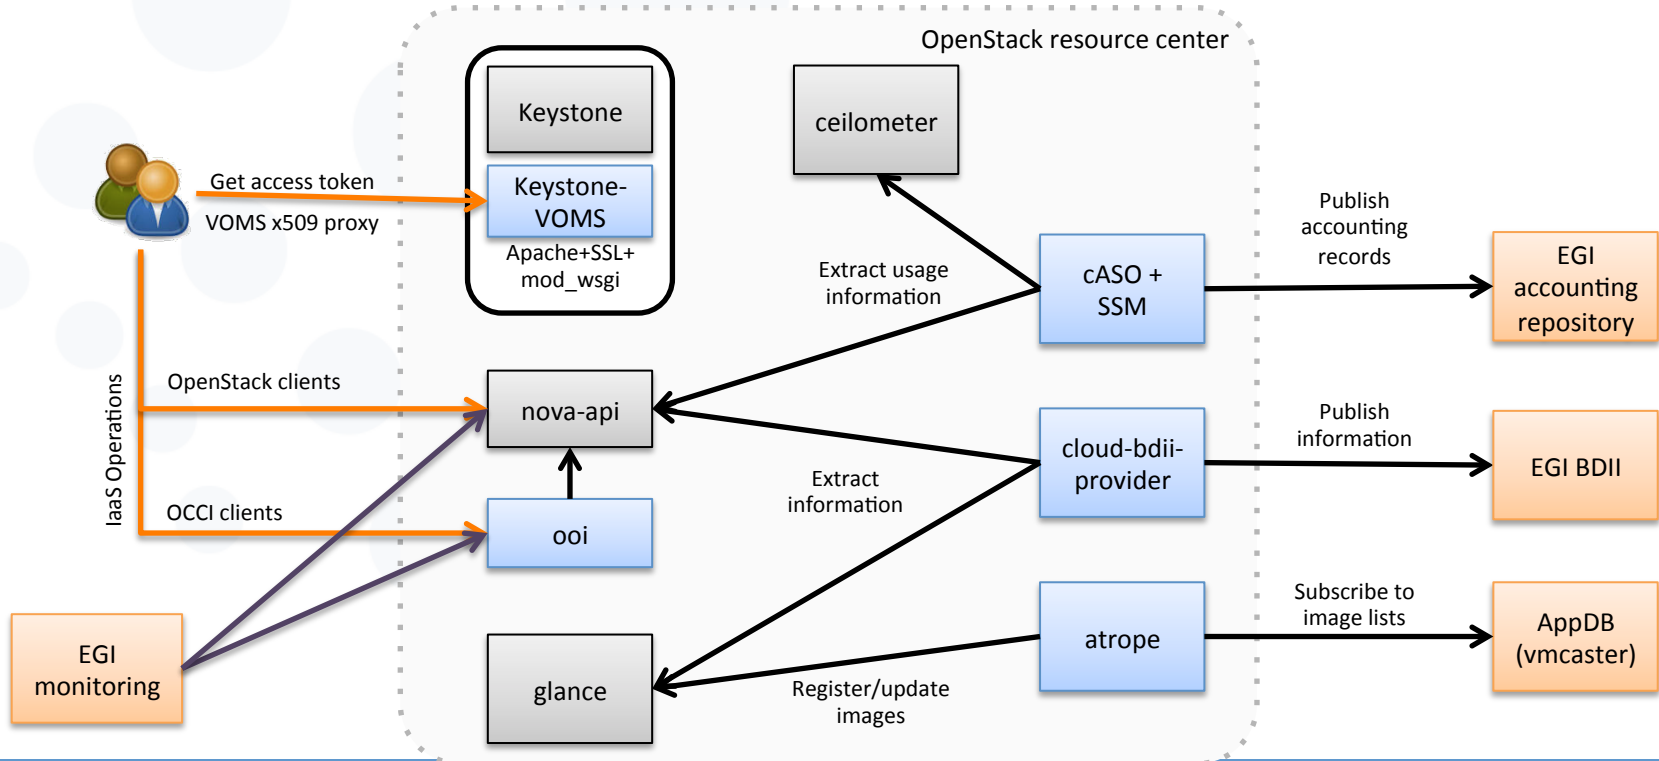
[Connected Items \(1\)](#)

OpenStack OCCI interface (ooi)

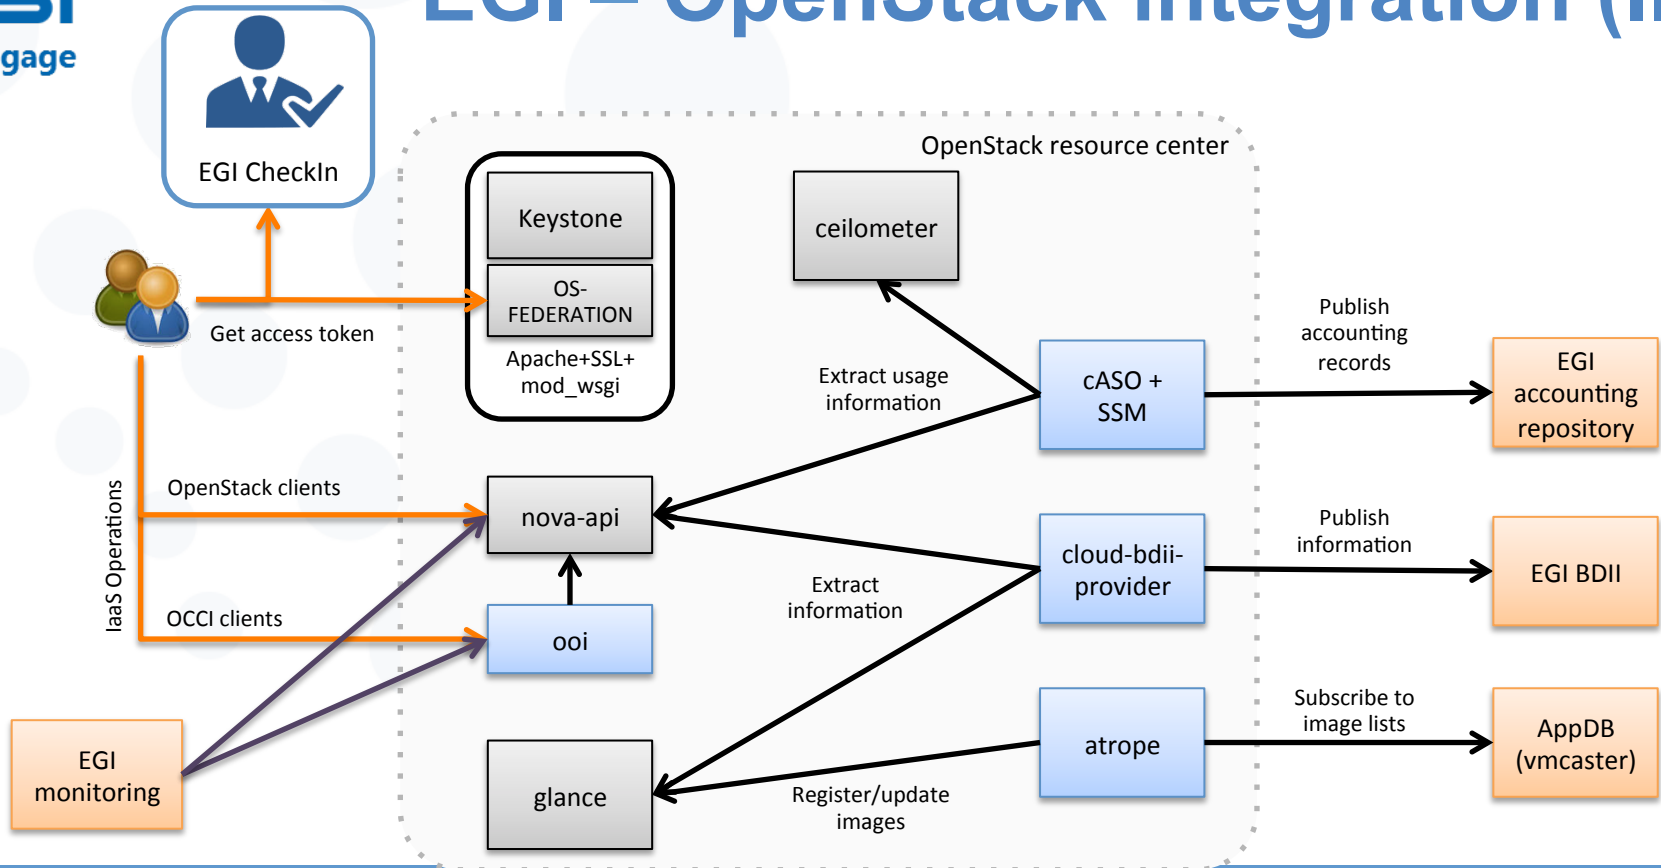
- **OCCI** (Open Cloud Computing Interface, OGF)
 - RESTful protocol and API focusing on cloud interoperability
 - Primarily for IaaS (manage VMs and Block Storage), extensible to other areas
- **ooi** (OpenStack OCCI interface)
 - Completely written from scratch OCCI implementation
 - Uses only public OpenStack APIs
 - Support for VM, volumes and network operations
 - Can be installed along an existing nova-api endpoint or as a separate WSGI application

<https://launchpad.net/ooi>

EGI – OpenStack integration (I)



EGI – OpenStack integration (II)



OpenStack FedCloud Appliance

- A single VM with all the components using the public OpenStack interfaces
 - Accounting, Information discovery, VMI replication
 - Packaged as Docker containers, available at docker hub
<https://hub.docker.com/u/egifedcloud/>
- Documentation:
 - https://wiki.egi.eu/wiki/MAN10#Integration_with_EGI_FedCloud_Appliance
- Appliance at AppDB:
 - <https://appdb.egi.eu/store/vappliance/fedcloud.integration.appliance.openstack>


Cloud Realms: build customized federations

Service	Cloud Realm Integration	EGI Technology
EGI Federated Service Management	mandatory	-
EGI Service Registry	mandatory	GOCDB
EGI AAI compliance	mandatory	Keystone-VOMS
EGI Accounting	mandatory , depends on EGI AAI compliance	cASO
EGI Monitoring	mandatory , done externally. Monitoring of IaaS interfaces (OCCI, OpenStack, CDMI) requires EGI AAI compliance	EGI ARGO
EGI Information Discovery	optional	cloud-bdii-provider
VMI replication	optional	atrope
VMI catalogue	optional	AppDB
IaaS standard interface	<i>optional</i> , all providers of the realm must provide homogeneous interface	ooi
EGI helpdesk	optional	GGUS


Communities using Cloud Compute (I)



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 flooding.




The National Bioinformatics Infrastructure of Sweden uses Cloud Compute to provide bioinformatics tools to their researchers, including high-profile 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

<https://www.egi.eu/use-cases/>

Communities using Cloud Compute (II)




Konrad Förstner and team used EGI Federacted cloud to analyse the RNAs produced by the Salmonella and humans at the same time, in the same experiments to discover that a piece of Salmonella RNA called PinT is heavily involved in what happens right after the infection. Results published in Nature.



VERCE platform is a science gateway, developed jointly by seismologists and IT experts, as a tool to create accurate earthquake simulations.

VERCE helped seismologists to make sense of the August 2016 Amatrice earthquake using the resources of EGI Cloud



Peachnote is a music score search engine and analysis platform. The system is the first of its kind and can be thought as an analog of Google Books Ngram Viewer and Google Books search for music scores. Peachnote is visited by tens of thousands of users every day from all over the world.

<https://www.egi.eu/use-cases/>

Pictures from wikimedia commons

What's coming

- Improve user experience
 - Certificate-less access with new EGI AAI
 - AppDB extension to provide VM management
- OCCI 1.2
 - Improved version of the standard, better networking support
- Go beyond IaaS
 - Exploit results from INDIGO-Project to offer PaaS to our users
- Get more involved in OpenStack
 - Scientific WG, identity federation, ...

Some references

- EGI: <http://www.egi.eu>
- Federated Cloud at EGI wiki:
[https://wiki.egi.eu/wiki/EGI Federated Cloud](https://wiki.egi.eu/wiki/EGI_Federated_Cloud)
 - Installation manual: <https://wiki.egi.eu/wiki/MAN10>
- EGI Federated Cloud list:
fedcloud-tf@mailman.egi.eu

Thank you for your attention.

Questions?

