The EGI Federated Cloud

Jérôme Pansanel

Journées France Grilles & LCG-France

IPNO – 06/18/2015

Outline

- The EGI Federated Cloud
- State-of-the-art and technologies
- Typical setups and use cases - through examples
- Plans, next steps
- Next steps to become a resource provider

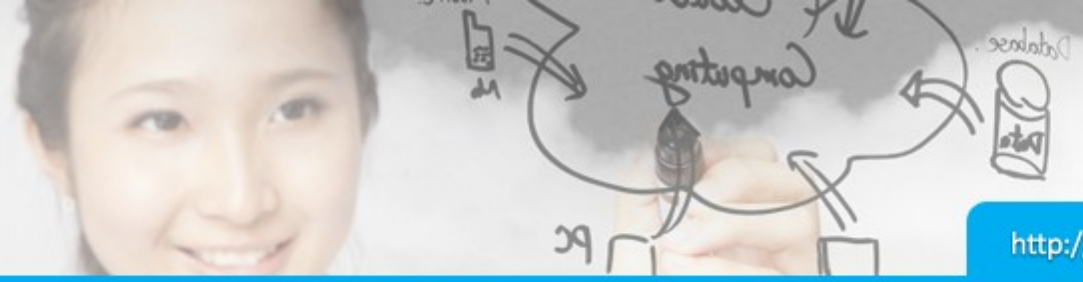

Outline

- **The EGI Federated Cloud**
- State-of-the-art and technologies
- Typical setups and use cases - through examples
- Plans, next steps
- Next steps to become a resource provider

The EGI Federated Cloud

The EGI Federated Cloud is federation of institutional private Clouds, offering Cloud Infrastructure as a Service to scientists in Europe and worldwide.

EGI Federated Cloud is based on:

- **Standards and validation**: federation is based on common **Open-Standards** – OCCI, CDMI, OVF, GLUE, etc...
- **Heterogeneous implementation**: no mandate on the cloud technology, the only condition is to expose the chosen interfaces and services.

The Environment

The EGI Federated Cloud is **providing access to digital resources on a flexible environment**.

Cloud resources	• Computing intensive experiments • Data intensive experiments • Hosting services
Standard based	• Heterogeneous implementation • OCCI, CDMI, GLUE2, OVF, etc.
Single system	• Set of independent cloud services • Uniform interfaces
Integrated in EGI infra	• AAI & monitoring • Accounting & information discovery • Service registry
VM Catalogue	• EGI Certified and Secure endorsed VMs • Register your VMs • Re-use public VMs
Target groups	• Individual researchers • Larger research communities or groups

Access to the EGI Federated Cloud

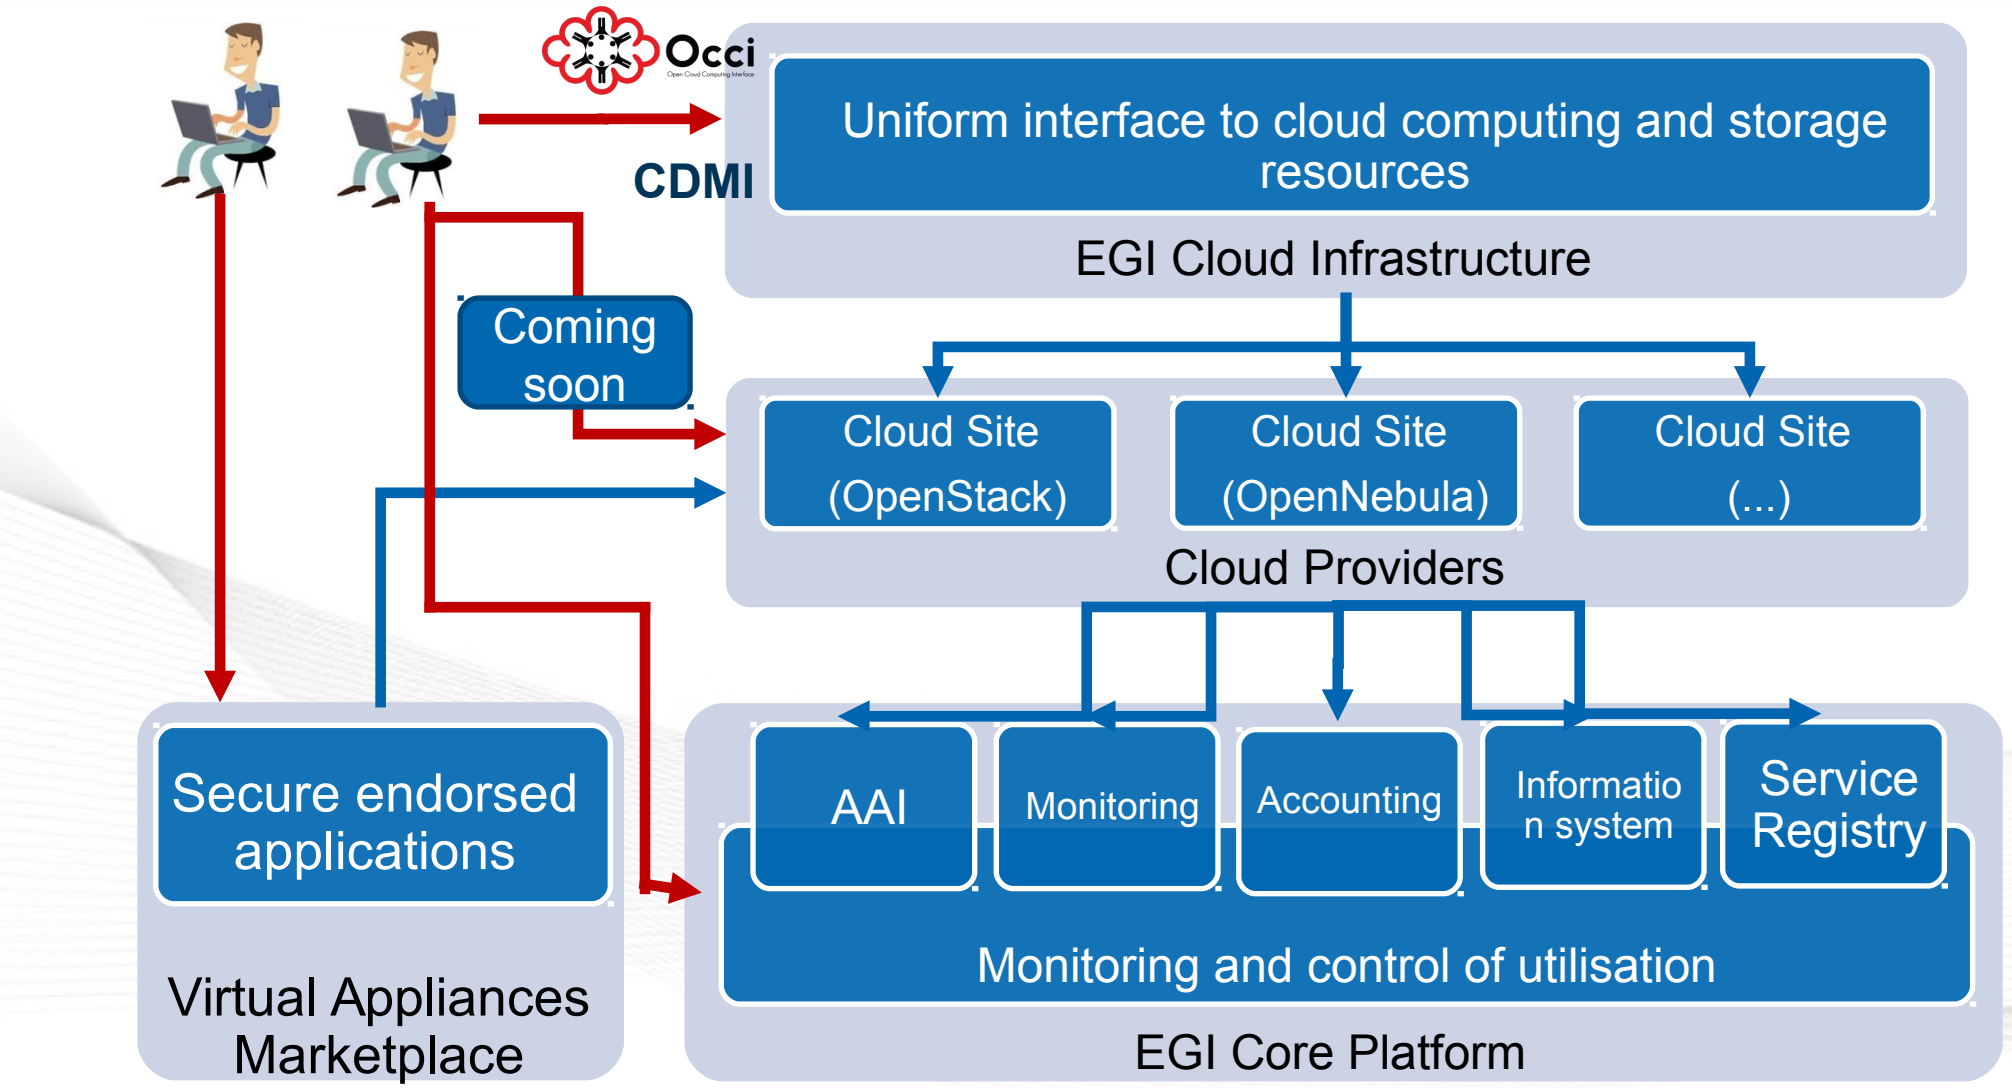

Resources

Resources

- **13 NGIs provide 22 certified resources**
- 4 NGIs currently integrating resources
- 5 NGIs with interested resource providers
- Worldwide interest & integration
  - Australia* (NeCTAR)
  - Africa* (SAGrid)
  - South Korea* (KISTI)
  - United States* (NIST, NSF)

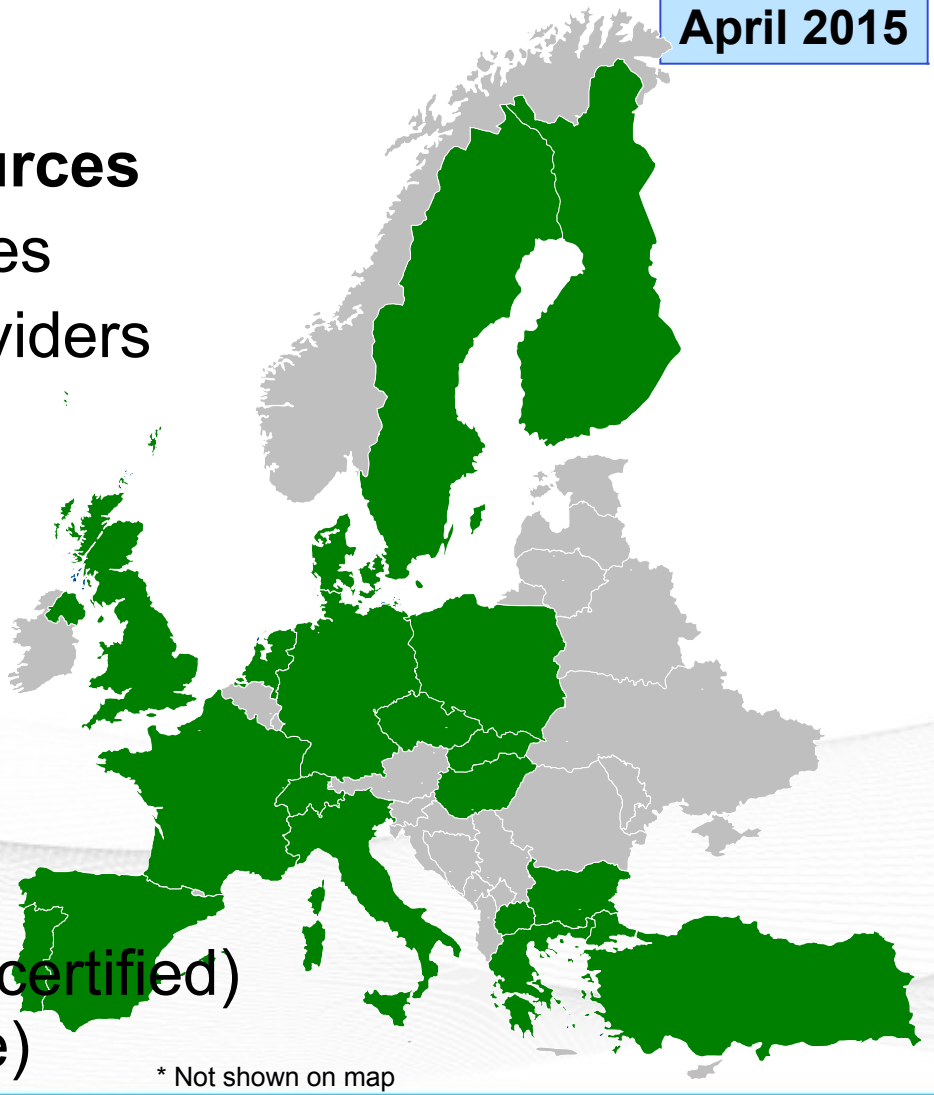

* Not shown on map

Usage – May/Dec 2014

- 244,913 (397,128) VMs certified (uncertified)
- 10.6M (12.3M) CPU hours (wall time)

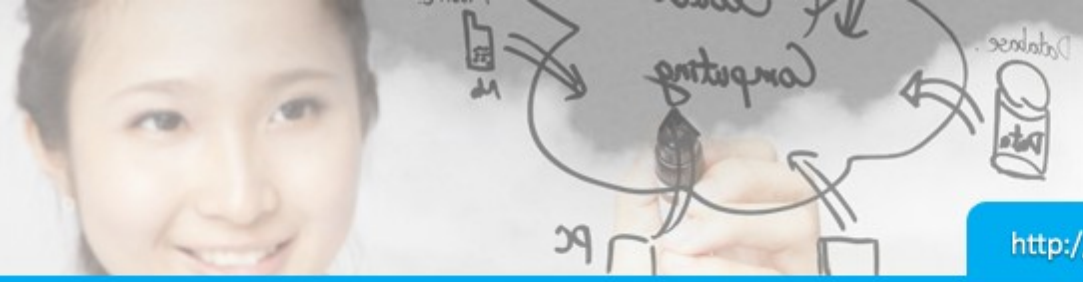

Outline

- The EGI Federated Cloud
- **State-of-the-art and technologies**
- Typical setups and use cases - through examples
- Plans, next steps
- Next steps to become a resource provider

Computing – VM Management

OCCI standard tools to easily manage VMs:

- Start, Suspend, Stop, Destroy, etc.

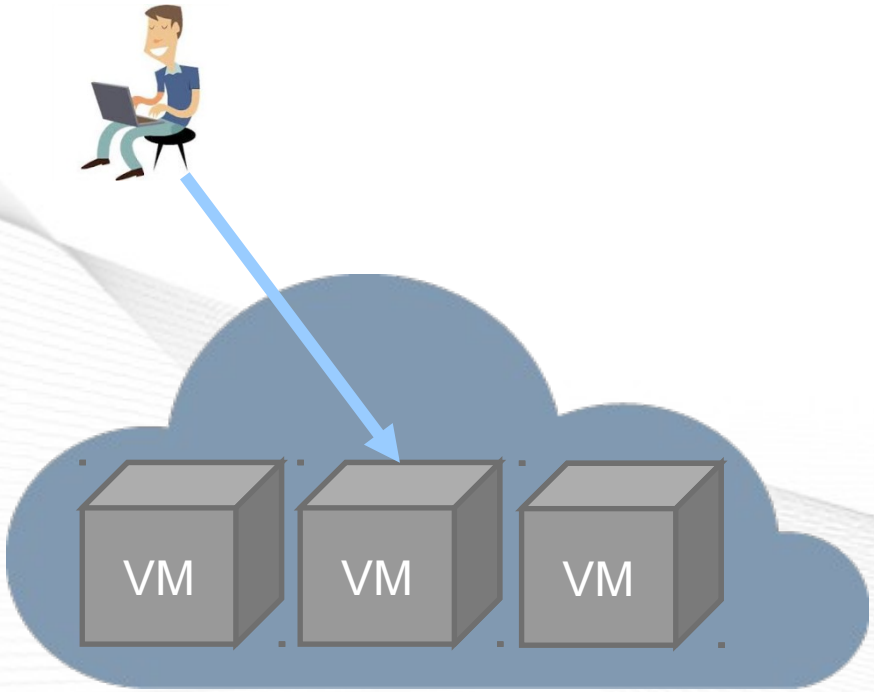

Simple usage

- OCCI standard
- rOCCI CLI client
- jOCCI – JAVA API

Scale to your needs

- Number of CPU cores
- Amount of RAM
- Local Disk

Contextualisation

- Configure your VM at startup
- Install SW at startup
- Complex deployment

Data Management

3 solutions available:

- VM Disk Space
- Block Storage (OCCI standard)
- Object Storage (CDMI standard)

VM Disk Space	• Disk stored in the VM image (<100 GB)
Block Storage	• Disk attached to a running VM (any size)
Object Storage	• External storage system • Access through REST API

Data Management – Block Storage

Persistent block level storage to use with VMs

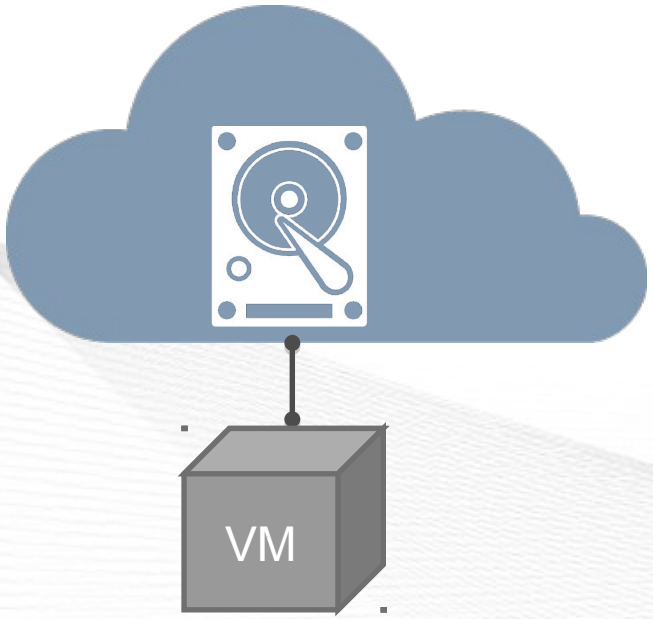

Simple usage	• Use as any other block device from VMs • Snapshotable
High Performance	• Consistent and low-latency performance • SSDs (in some sites)
Scale to your needs	• From GB to TB • Create and attach to VMs on demand

Data Management – Object Storage

Data storage infrastructure for storing and retrieving data from anywhere at any time

API Access	• Simple REST APIs for managing and accessing data
Scalable	• Store as much data as needed. • Get accounted only for the space used.
Sharing	• Define ACLs on each object, share publicly your data

Appliances Marketplace

Web-based VM image catalogue with automatic deployment on the EGI Federated Cloud sites

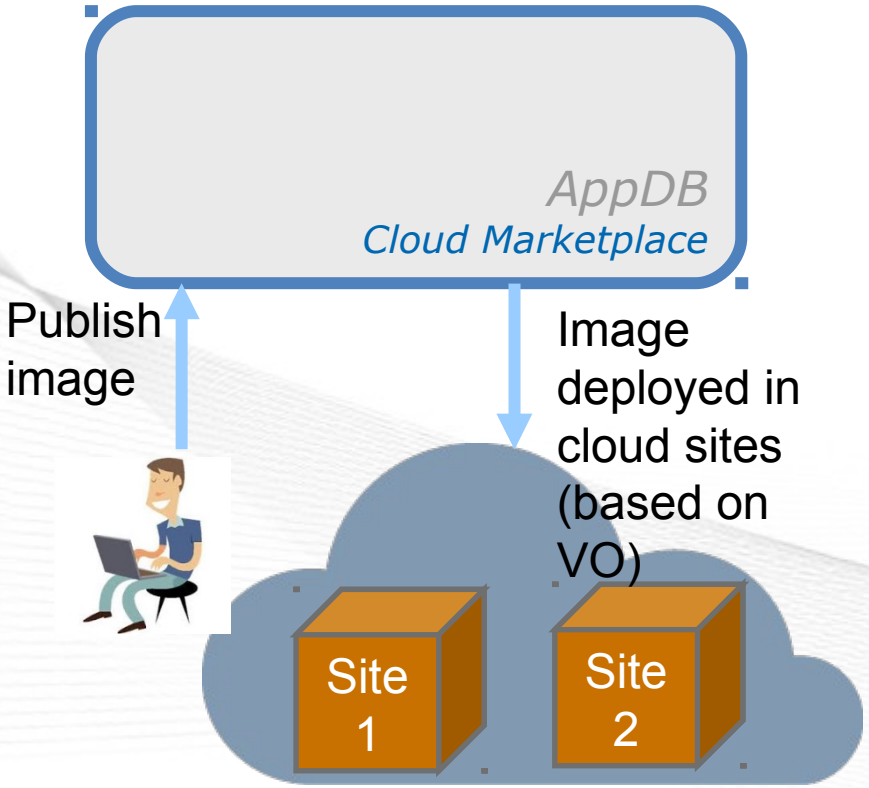

Cloud VM image repository	• Publish your image • Re-use images certified by EGI • Manage versions
Manage your VM images	• Sharing • Create private images
Automatic deployment on the cloud sites	• Images automatically deployed on the cloud sites • Automatic update

https://appdb.egi.eu/store/vappliance/

Security

Secure endorsed application and service deployment

EGI CSIRT
Computer Security Incident Response Team

https://wiki.egi.eu/wiki/EGI_CSIRT:Main_Page

EGI VM Images	• Endorsed by the EGI security team • Periodically updated
User VM Images	• Guidelines to create secure images • Each user can ask for an EGI endorsement
Security monitoring	• Continuous monitoring of the running VMs • Alarms and immediate actions in case of security issues

Information Discovery (BDII)

Information about the available resources and their current status

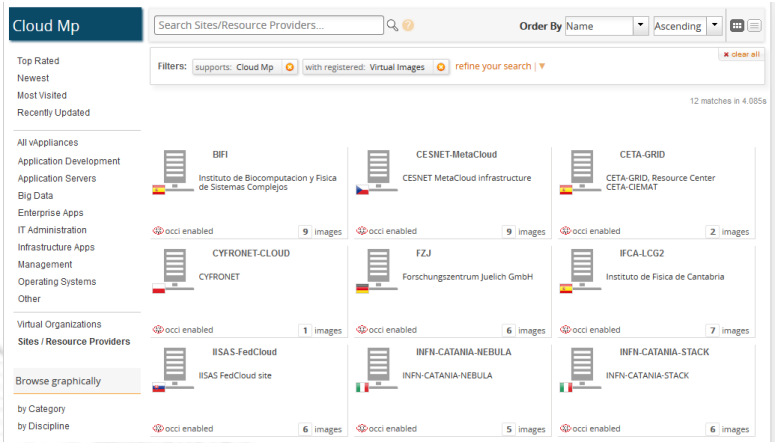

Sites
- List of the certified EGI Federated Cloud sites

VM images
- List of images available on each site
- Get IDs to be used with the rOCCI client

Views
- High level view for final users (VA marketplace)
- Advanced view for expert users through LDAP

Accounting

Data about usage of the FedCloud resources

The following table shows the distribution of Total number of VM run grouped by SITE and DATE.

Total number of VM run by SITE and DATE

SITE	Nov 2014	Dec 2014	Jan 2015	Feb 2015	Mar 2015	Apr 2015	Total	%
100IT	716	709	748	675	743	48	3,639	1.08%
BIFI	11,113	15,293	18,209	42,974	47,682	0	135,271	40.00%
CERN-PROD	13,310	185	437	8,085	171	34	22,222	6.57%
CESGA	1,726	929	0	1	556	52	3,264	0.97%
CESNET-MetaCloud	1,258	1,407	792	837	770	49	5,113	1.51%
CETA-GRID	153	99	2,307	227	454	45	3,285	0.97%
CYFRONET-CLOUD	456	0	0	235	751	47	1,489	0.44%
FZJ	5,530	17,979	10,385	21,145	3,088	47	58,174	17.20%
GoeGrid	1,504	798	751	681	752	49	4,535	1.34%
HG-09-Okeanos-Cloud	493	229	241	185	179	0	1,327	0.39%
IFCA-LCG2	4,426	650	621	539	537	35	6,808	2.01%
IISAS-FedCloud	2,970	6,014	5,357	2,766	3,435	207	20,749	6.13%
IN2P3-IRES	0	0	0	188	729	44	961	0.28%
INFN-CATANIA-NEBULA	663	705	779	672	767	49	3,635	1.07%
INFN-CATANIA-STACK	115	202	847	697	774	47	2,682	0.79%
INFN-PADOVA-STACK	5,482	8,652	13,131	11,162	7,683	1,075	47,185	13.95%
MK-04-FINKICLOUD	284	344	749	688	227	0	2,292	0.68%
NCG-INGRID-PT	667	674	712	593	714	30	3,390	1.00%
PRISMA-INFN-BARI	2,460	753	791	208	212	1	4,425	1.31%
SZTAKI	551	572	487	707	31	0	2,348	0.69%
TR-FC1-ULAKBIM	157	147	752	707	499	28	2,290	0.68%
UPV-GRYCAP	756	574	257	888	600	49	3,124	0.92%
Total	54,790	56,915	58,353	94,860	71,354	1,936	338,208	
Percentage	16.20%	16.83%	17.25%	28.05%	21.10%	0.57%		

VO Group/Role: fedcloud.egi.eu - ALL Groups and Roles
ROC/Country: ALL ROCs
Infrastructure: Cloud
Order by: Number of VMs — Total number of VMs run
Period: Start year: 2014 Start month: 3 End year: 2015 End month: 2

Refresh

VO (ALL Groups and Roles) Users information

Mar 2014 - Feb 2015

The following table shows the Usage of the Users ordered by Total number of jobs and the Total Usage of the Other Users. A detailed view can be obtained by selecting an individual user

VO Users ordered by Total number of jobs

#	User ID	Jobs #	Jobs %	CPU time Hrs	CPU time %	Norm. CPU time Hrs	Norm. CPU time %	WCT Hrs	WCT %	Norm. WCT Hrs	Norm. WCT %	Monetary Cost €	Monetary Cost %	CPU Efficiency %	Avg. CPU time Hrs	Avg. WCT Hrs
1	User 1	34,158	76.8%	102,855,939	5.4%	0	0.0%	4,359,851,264	64.5%	0	0.0%	166,977,133	15.5%	2.4	3011.18	127637.78
2	User 2	2,040	4.6%	79,340,745	4.2%	0	0.0%	79,340,745	1.2%	0	0.0%	80,398,622	7.4%	100.0	38892.52	38892.52
3	User 3	1,826	4.1%	2,154,330	0.1%	0	0.0%	2,154,330	0.0%	0	0.0%	0	0.0%	100.0	1179.81	1179.81
4	User 4	1,138	2.6%	324,439,069	17.1%	0	0.0%	324,439,069	4.8%	0	0.0%	258,438,790	23.9%	100.0	285095.84	285095.84
5	User 5	545	1.2%	64,898,295	3.4%	0	0.0%	137,729,788	2.0%	0	0.0%	54,738,970	5.1%	47.1	119079.44	252715.20
6	User 6	539	1.2%	126,928,755	6.7%	0	0.0%	146,890,024	2.2%	0	0.0%	998,313	0.1%	86.4	235489.34	272523.24
7	User 7	403	0.9%	151,702,037	8.0%	0	0.0%	268,360,390	4.0%	0	0.0%	105,699,879	9.8%	56.5	376431.85	665906.67
8	User 8	293	0.7%	22,456,541	1.2%	0	0.0%	22,456,541	0.3%	0	0.0%	22,273,052	2.1%	100.0	76643.48	76643.48
9	User 9	227	0.5%	122,210,568	6.4%	0	0.0%	167,553,207	2.5%	0	0.0%	95,657,093	8.9%	72.9	538372.54	738119.85
10	User 10	193	0.4%	28,527,866	1.5%	0	0.0%	116,496,393	1.7%	0	0.0%	24,203,168	2.2%	24.4	147812.78	603608.25
11	User 11	177	0.4%	2,859,069	0.2%	0	0.0%	2,859,069	0.0%	0	0.0%	2,897,190	0.3%	100.0	16152.93	16152.93
12	User 12	120	0.3%	8,908,708	0.5%	0	0.0%	9,296,073	0.1%	0	0.0%	2,742,557	0.3%	95.8	74239.23	77467.28
13	User 13	113	0.3%	4,134,284	0.2%	0	0.0%	4,150,007	0.1%	0	0.0%	538	0.0%	99.6	36586.58	36725.73

- **Groups (VO) and users**
  - N. of VMs created for each group (VO) or user
  - Total CPU times
  - Total RAM, etc.
- **Sites**
  - N. of VMs
  - Total CPU times
  - Total RAM
  - Disk size, etc.
- **VMs**
  - N. of cores
  - RAM
  - Disk size
  - VM image, etc.

Outline

- The EGI Federated Cloud
- State-of-the-art and technologies
- **Typical setups and use cases - through examples**
- Plans, next steps
- Next steps to become a resource provider

Use Case Models

The EGI Federated Cloud is able to support different use case models

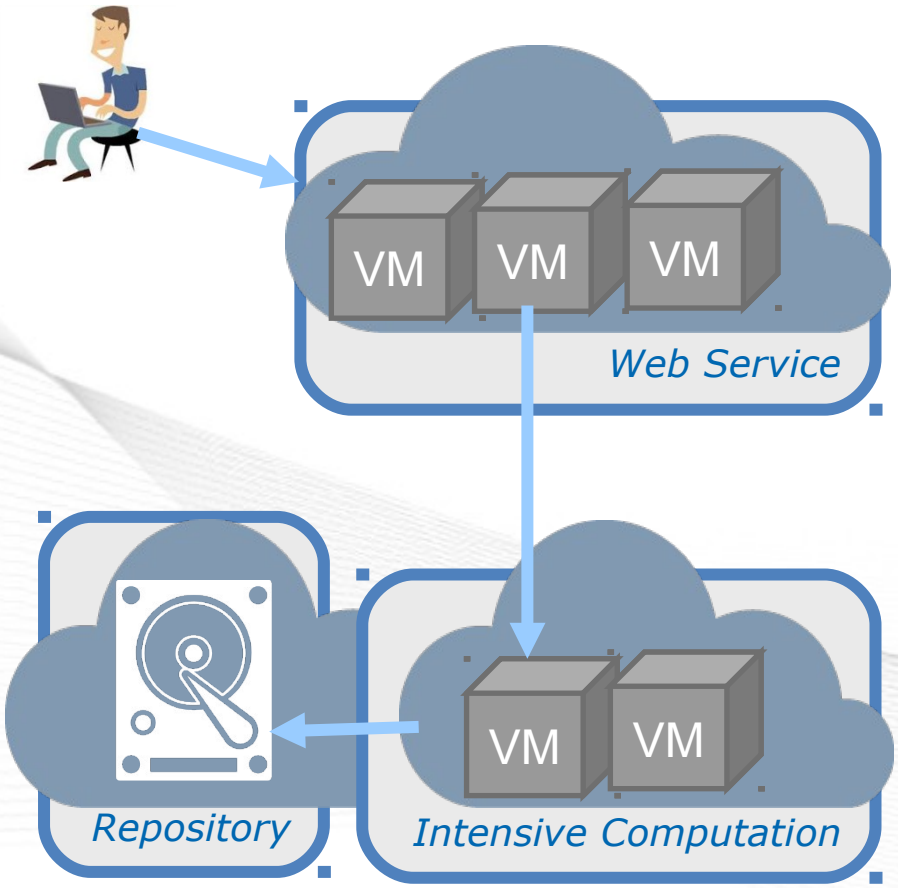

You can combine them!

Compute & Data intensive	• Heavy computation • Large Memory
Datasets Repository	• Store & manage large datasets
Web Services	• Hosting web-sites, web services, portal, etc. • Exploit cloud features as elasticity

Example: Analysis of RNA-Seq Data

Pipeline for the computational evaluation of RNA-Seq. data

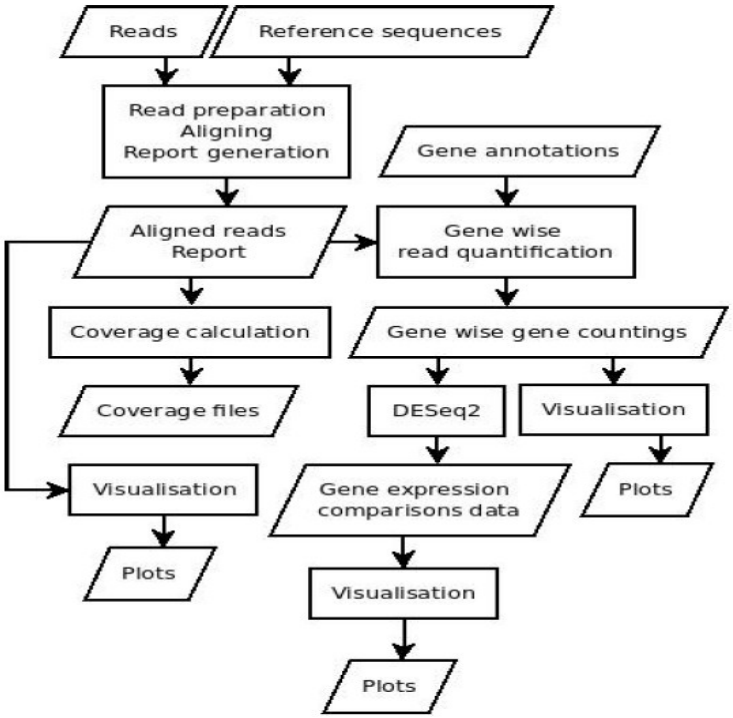

Source: Konrad U. Förstner

UNIVERSITÄT WÜRZBURG

Usage Model	• **Heavy computation** • **Large Memory**
Scientific Disciplines	• **Bioinformatics**
Deployment in the FedCloud	• **VMs with 24 cores, 128 GB of RAM** • **Block storage up to 3 TB**

Outline

- The EGI Federated Cloud
- State-of-the-art and technologies
- Typical setups and use cases - through examples
- **Plans, next steps**
- Next steps to become a resource provider

Next Steps

The EGI Federated Cloud Task Force continue to work to evolve and improve the Cloud Infrastructure

Compute	• OCCI 1.2 • MS Azure • Native CMF interfaces
High level tools	• More PaaS and SaaS integrated in the FedCloud
Storage	• OCCI advanced for Block Storage • CDMI for OpenStack • Native CMF interfaces
Network	• Advanced network management

References

EGI Federated Cloud resources

- Talk: https://www.egi.eu/news-and-media/presentations/#fedcloud
- Wiki site: http://go.egi.eu/fedcloud
- User support: https://wiki.egi.eu/wiki/Federated_Cloud_user_support
- User support e-mail: support@egi.eu
- Federated Cloud Communities: https://wiki.egi.eu/wiki/Federated_Cloud_Communities

Related Standards:

- OCCI: http://occi-wg.org
- CDMI: http://cdmi.sniacloud.com/

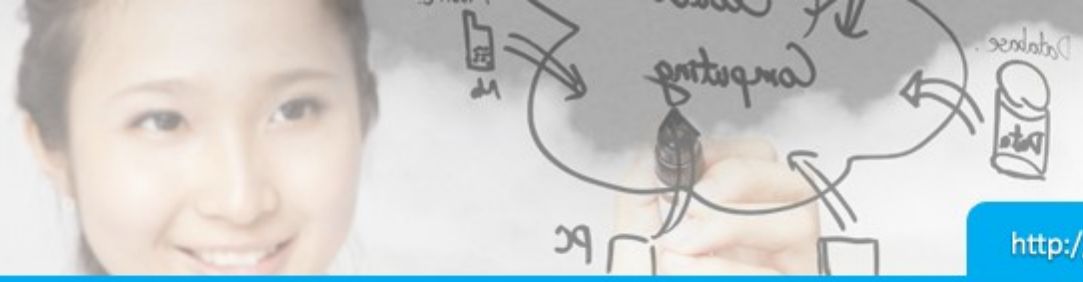

Outline

- The EGI Federated Cloud
- State-of-the-art and technologies
- Typical setups and use cases - through examples
- Plans, next steps
- **Next steps to become a resource provider**

Before to start

https://wiki.egi.eu/wiki/MAN10

EGI FedCloud supports 3 middlewares

- OpenStack
- OpenNebula
- Synnefo

Prerequisities & Limitations

- Hardware (minimal hw requirements, ...)
- DNS names, X.509 certificates
- Register in fedcloud VO
- Registration in AppDB to have access to private EGI VM image repository
- What operating systems are supported

Installation

https://wiki.egi.eu/wiki/MAN10

- Install the Cloud middleware (i.e. OpenStack)
  - Icehouse, Juno
  - Neutron with per-tenant routers and private networks
  - Swift if providing CDMI
- Valid IGTF-trusted host certificates for Keystone / https auth
- Keystone-VOMS plugin
- Configure the supported VOs (voms.json, vomses file, tenant, networks, accounting, image list)
- OCCI plugin
- Accounting extractor based on cASO
- Information system (BDII provider plugin)
- EGI Image Management (vmcatcher, glancepush)
- Registration of the service in the GOCDB

Test your Infrastructure

https://wiki.egi.eu/wiki/MAN10

- Verify on Cloudmon your services:
  https://cloudmon.egi.eu/nagios
- Check that your site is publishing Cloud information in the BDII
- Check that the supported images are listed in the BDII
- Execute the site manual certification tests:
  https://wiki.egi.eu/wiki/HOWTO04_Site_Certification_Manual_tests#Check_the_functionality_of_the_cloud_elements
- Ask to join the FedCloud as certify Resource Provider
- Check the accounting portail about your site