EGI FedCloud; Connecting Researchers to Clouds

Dr David Wallom

Chair EGI Federated Cloud Task Force

Associate Director, Oxford eResearch Centre, University of Oxford

david.wallom@oerc.ox.ac.uk

Outline

- EGI, new challenges and cloud computing.
- Learning from other projects and activities.
- Federation model objectives, deliverables, mandate and membership.
- Standards and core infrastructure services.
- Federation test bed.
- Use cases.
- Conclusions.

Operational Infrastructure

Registered Users: 21714 VOs: 233
LCPUs: 470,000 Disk: 143PB Tape: 138PB
Jobs: 1.62 million/day

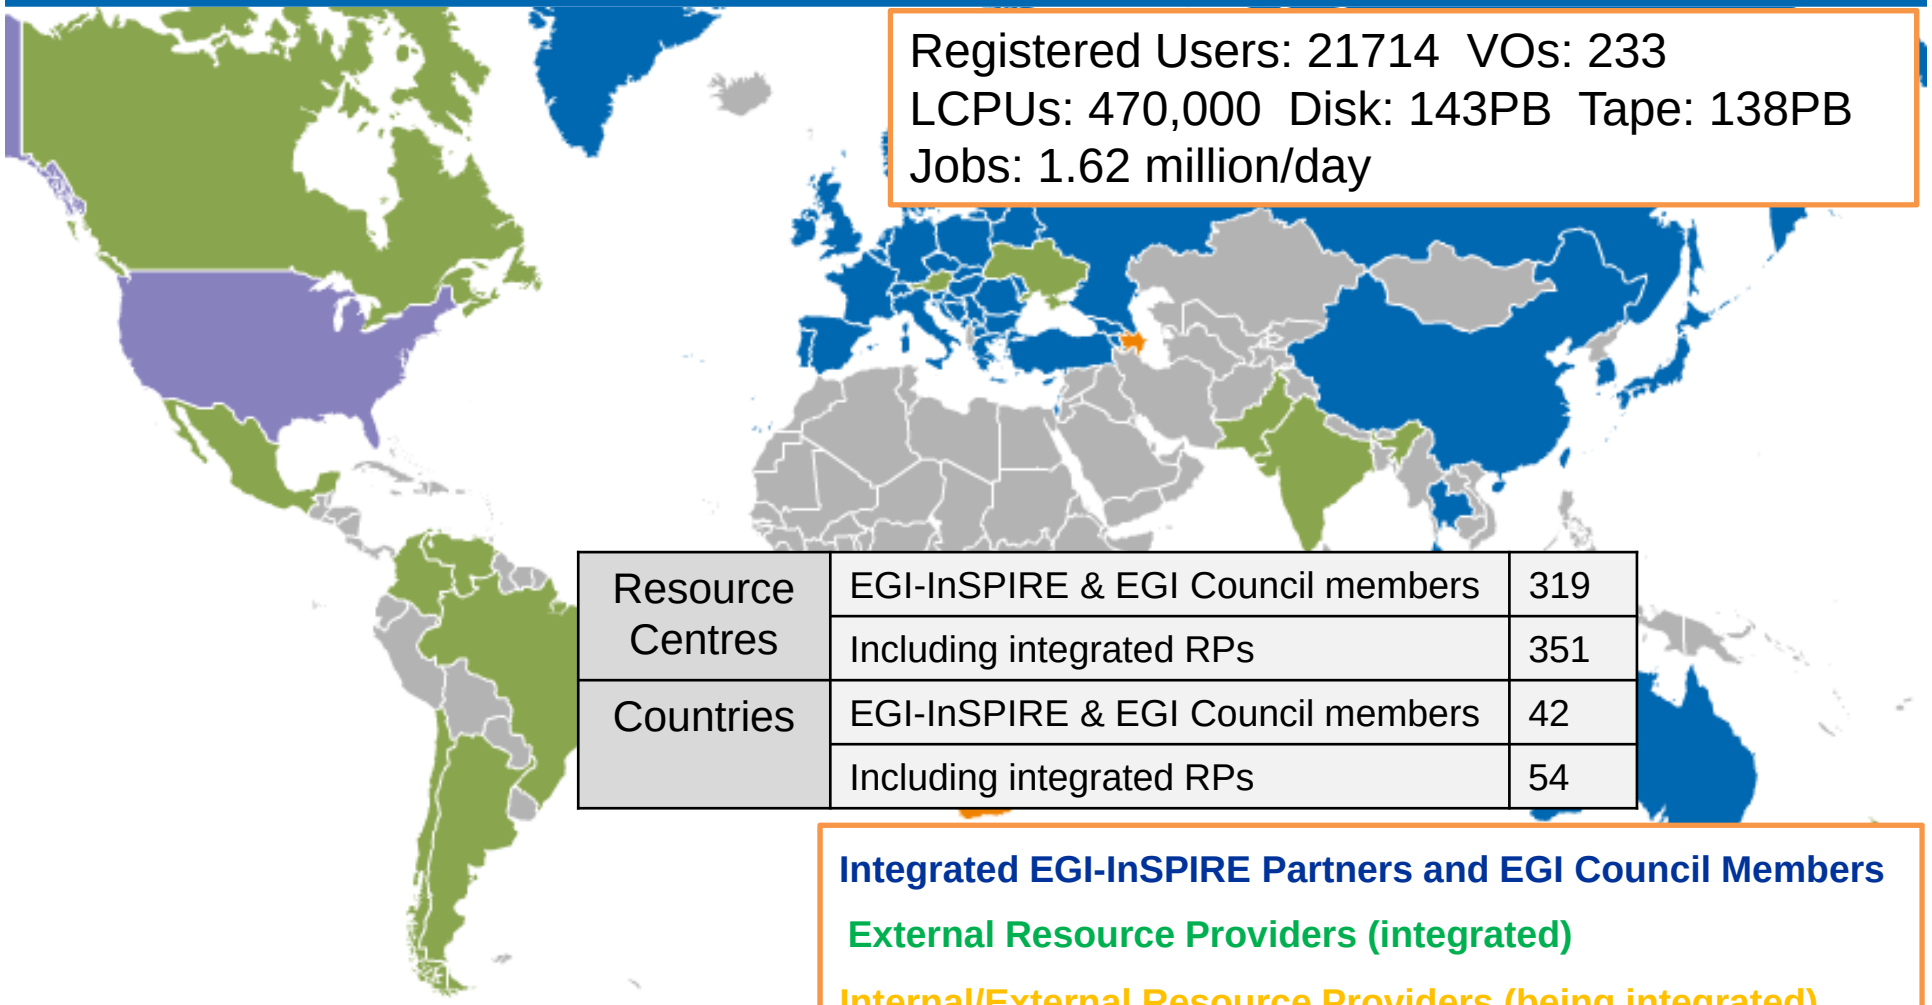

Resource Centres	EGI-InSPIRE & EGI Council members	319
	Including integrated RPs	351
Countries	EGI-InSPIRE & EGI Council members	42
	Including integrated RPs	54

Integrated EGI-InSPIRE Partners and EGI Council Members

External Resource Providers (integrated)

Internal/External Resource Providers (being integrated)

Peer Resource Providers

Thanks to Steven Newhouse

EGI New Challenges and Cloud Computing

- New communities and application design models
- Need for long running services (not only jobs).
- Workflows that integrate local and remote systems.
- Integrating community-specific resources (sensors, antennas, repositories, ...).
- Since ~2010 the trend has been for resource providers to move Grid middleware in virtualised environments.

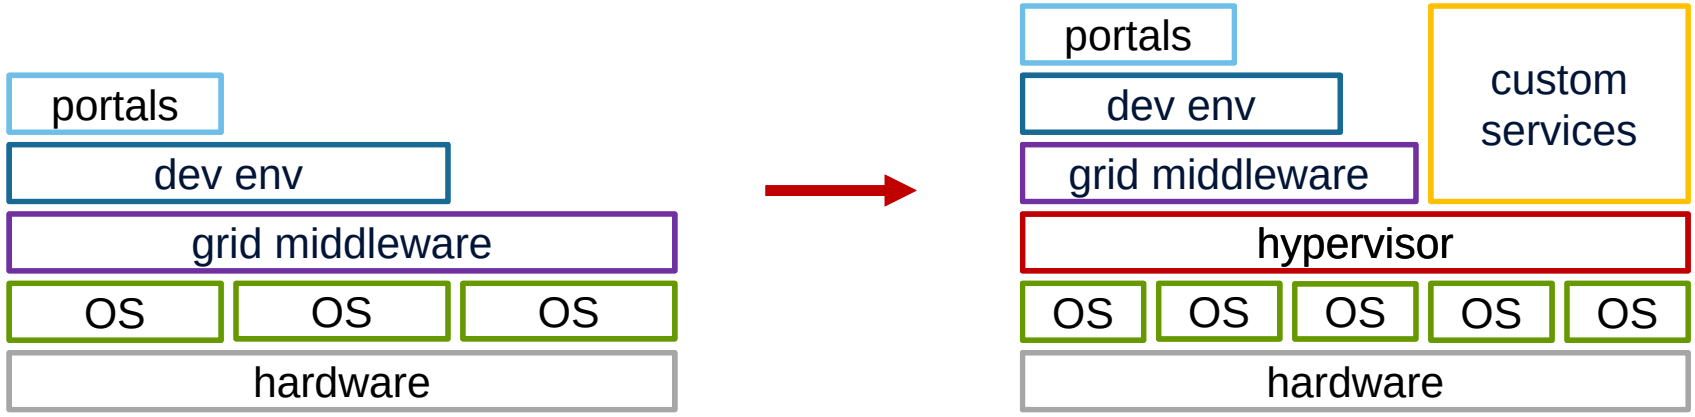

Flexible Services for the Support of Research (FleSSR) – 2008 -> 2011

6 Partners

- Academic and industrial;
- 3 cloud infrastructures.

Goals

Building federated cloud infrastructure, extending the use of UK NGS central services with cloud brokering and accounting.

Use cases

- Multi Platform Software Development;
- On demand Research data storage.

FleSSR Infrastructure assumptions & Architecture

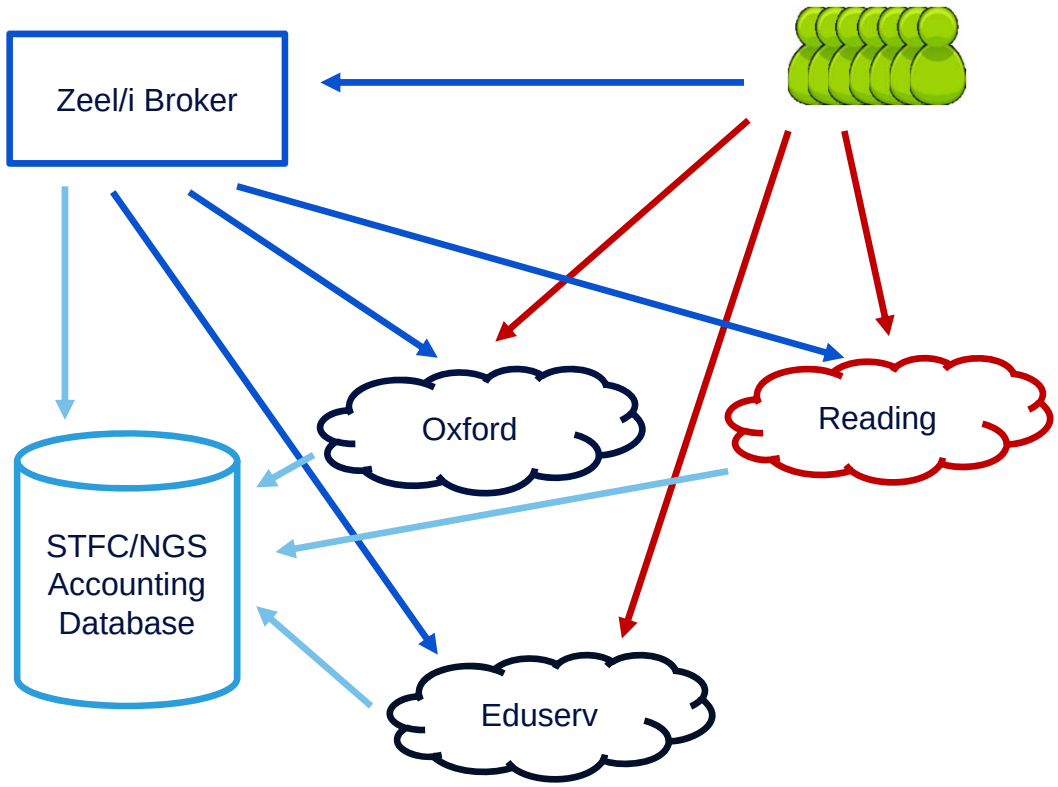

- **Local/Global**: services depends either on local or global access. Cloud brokering is not mandatory for AWS-like service access;
- **Multiple identities**: every user may have multiple identities, both local and global;
- **Only personal identities**: group identities are not implemented. The management of every single identity is left to the legally responsible user;
- **Multiple AA technologies**: AA may differ depending on local and global policies/technologies;
- **Multiple accounting**: every single identity is accounted for its usage. Every individual may get multiple invoices.

FleSSR Use Cases

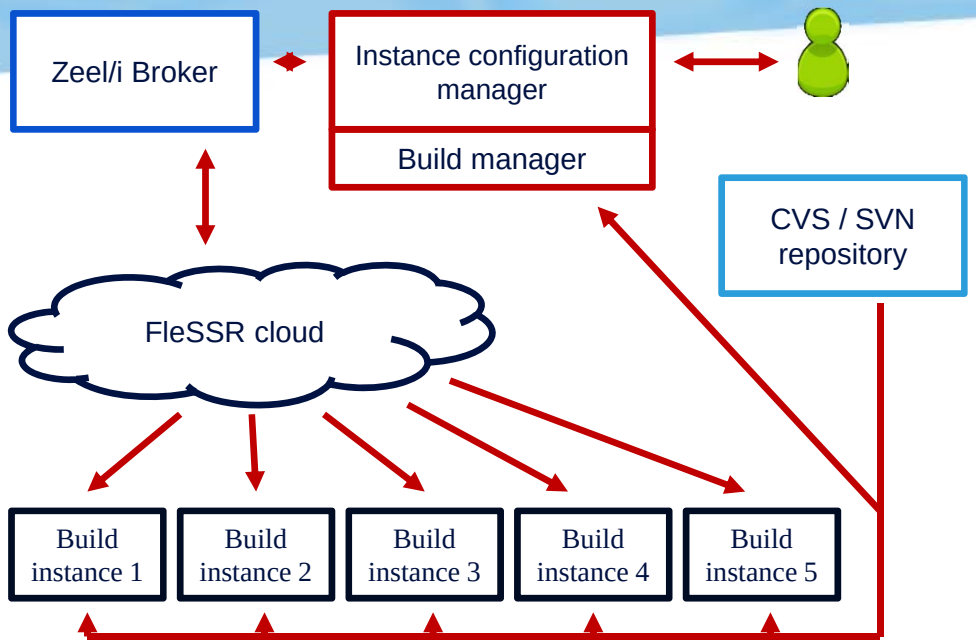

Multi Platform Software Development

- Management of multiple VM instances

On demand Research data Storage

- Management of multiple storage volumes

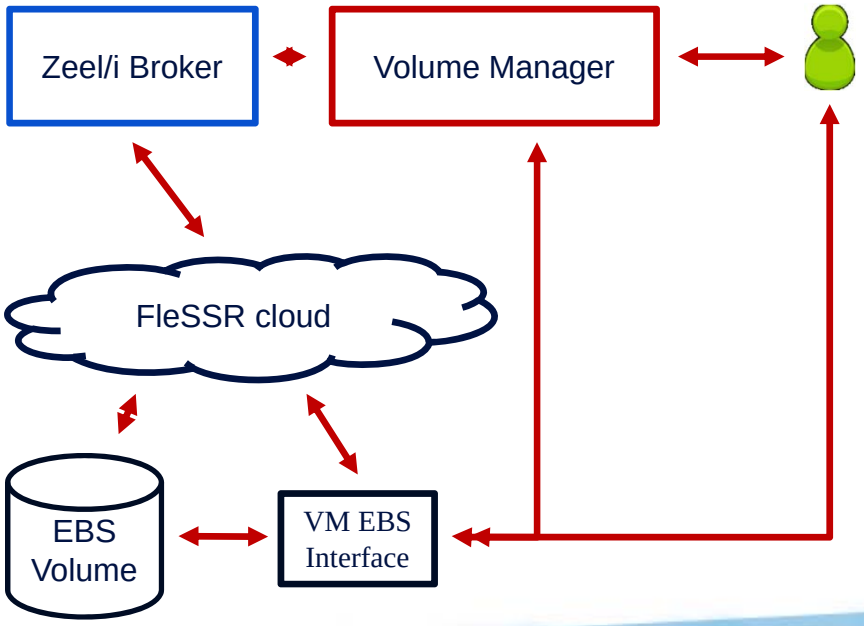

EGI Federation Model

- **Standards and validation**: emerging standards for the interfaces and images – OCCI, CDMI, OVF.
- **Resource integration**: Cloud Computing to be integrated into the existing production infrastructure.
- **Heterogeneous implementation**: no mandate on the IaaS cloud technology.
- **Provider agnosticism**: the only condition to federate resources is to expose the chosen interfaces and services.

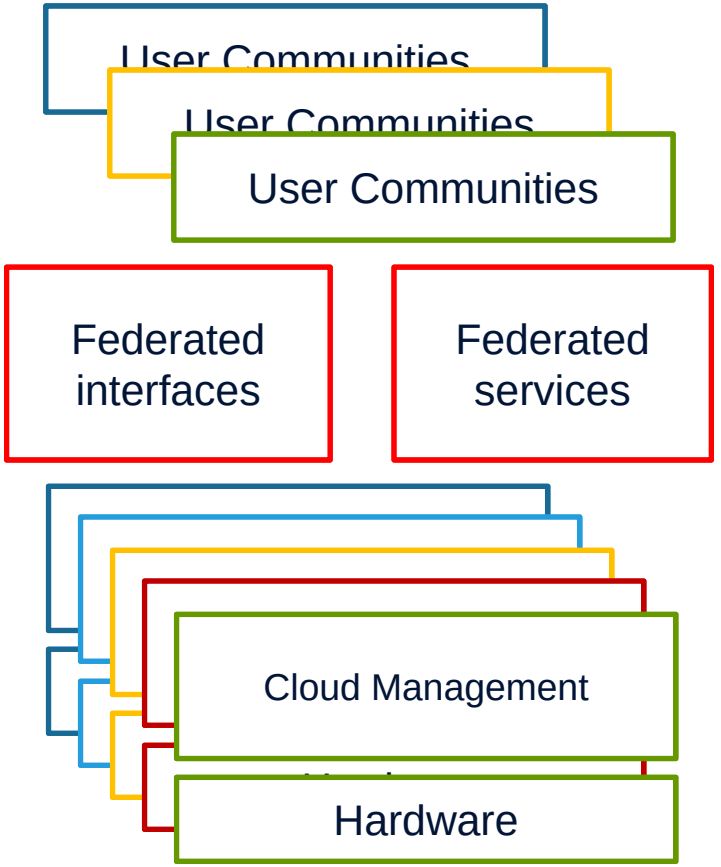

EGI Task Force Federated Clouds

- Stakeholders
  1. Resource Centres *offer Cloud services*
  2. Technology Providers *tender solutions*
  3. User communities *Consume cloud services*

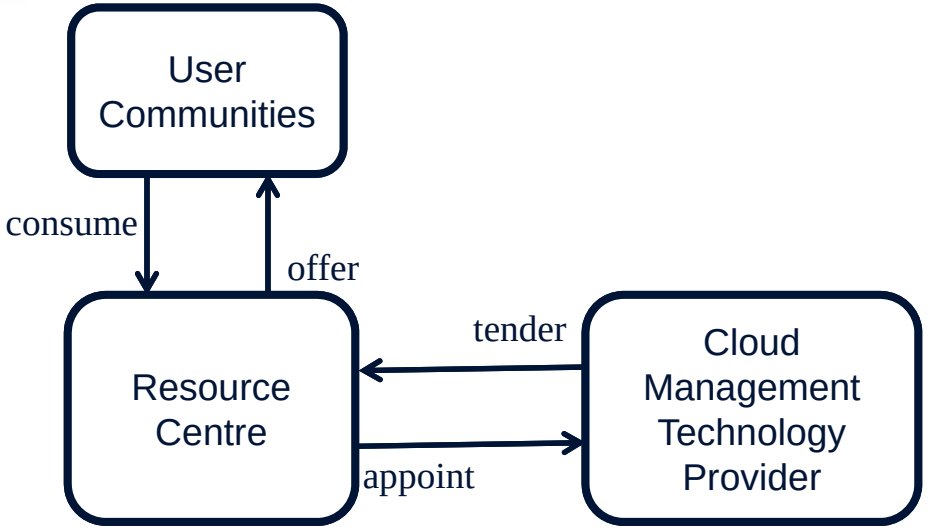

- 10 scenarios = 10 key *federated* cloud capabilities

1. Manage VM instances
2. Data access/transfer interface
3. Cloud service information federation
4. Resource consumption management
5. Cloud service availability
6. Notification & Automation
7. Federated AAI
8. VM Image Management
9. Brokering
10. Contextualisation

PreProduction – September 2013

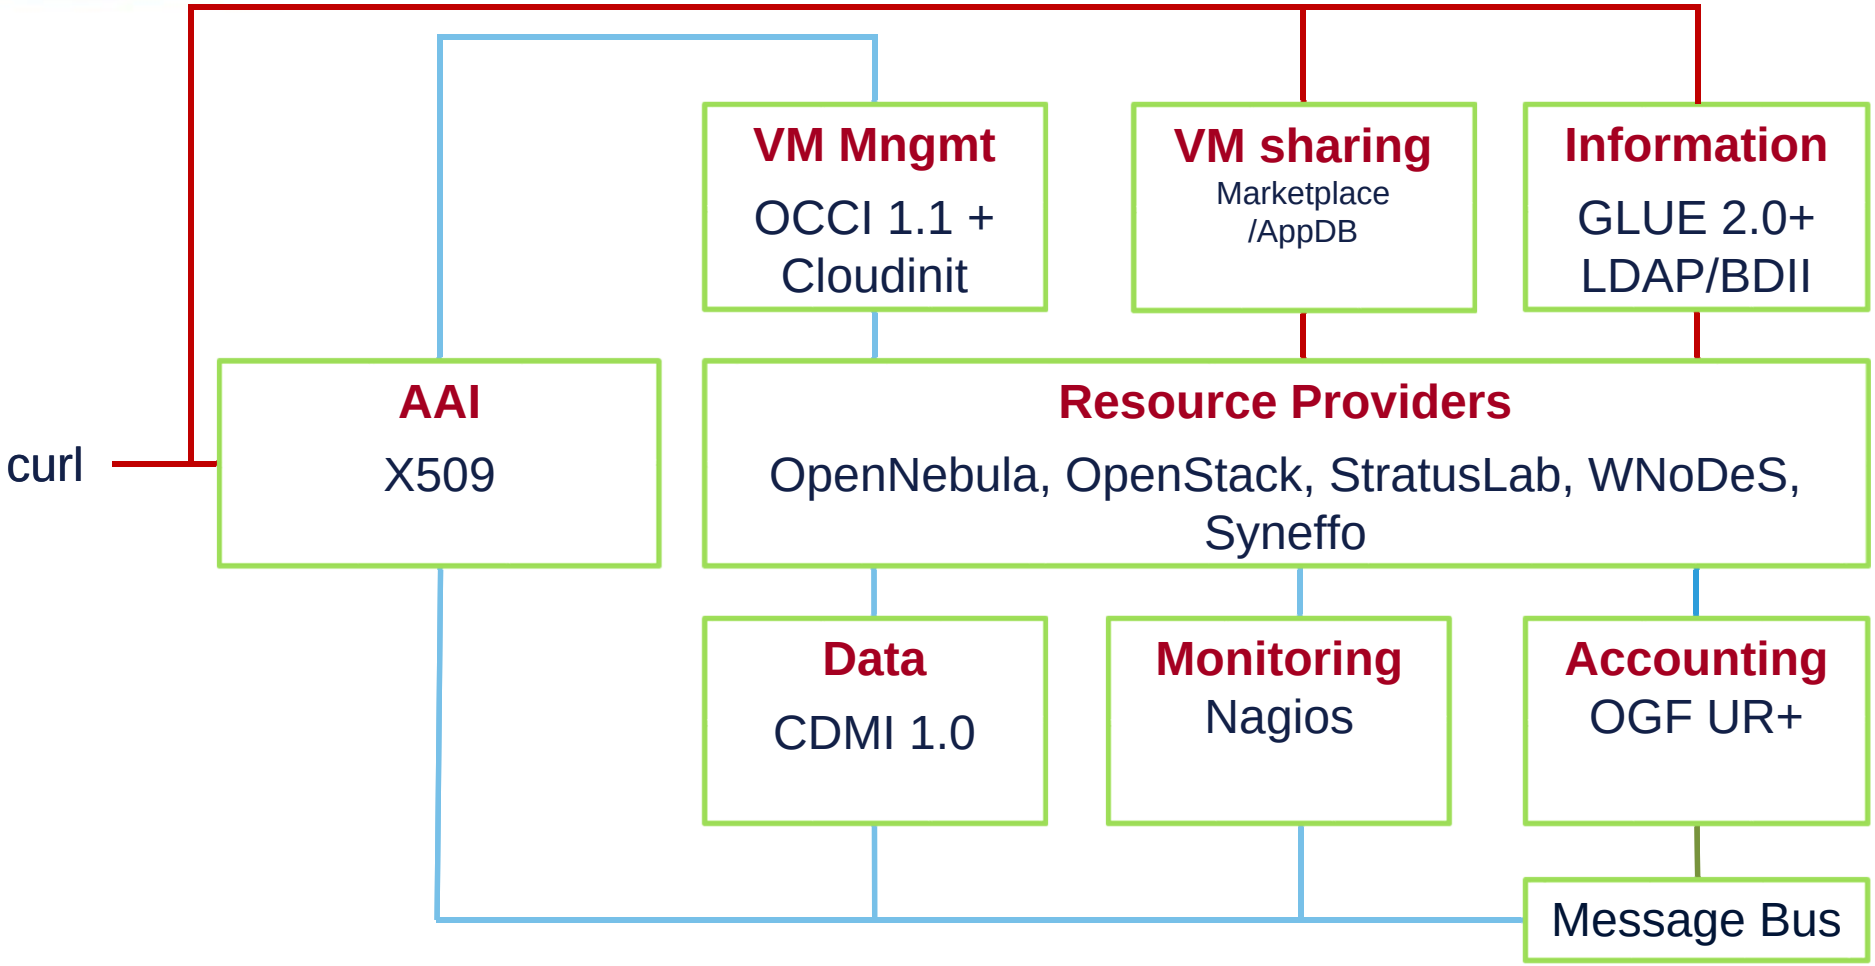

Federation testbed technologies and services

EGI Core Infrastructure

Service Availability
SAM

Configuration DB
GOCDB

Accounting
APEL

External Services

Certification Authorities

ID Management Federations

AAI

VOMS Proxy

User Interfaces

OCCI Clients
rOCCI; WNoDeS-CLI
SlipStream, VMDIRAC,
CompatibleOne

CDMI Clients
Libcdmi-java

Resource Providers

VOMS proxy

OCCI server

CDMI server

OpenNebula
OpenStack
StratusLab
WNoDeS
Syneffo,
others

Marketplace client

SSM

Vmcatcher
Vmcaster

LDAP

FedCloud Core Services

Image Metadata
Marketplace

Appliance Repository

Information System
Top-DBII

Monitoring
Nagios + community probes

VO
PERUN server

Objectives and Deliverables

- **Engagement**: identify and work with resources providers, technology providers, and user communities.
- **Integration**: integration of cloud resources within EGI's production infrastructure – e.g. monitoring, accounting and information publishing.
- **Recommendations**: identify issues that need to be addressed by other areas of EGI – e.g. policies, operations, support and dissemination.

- **Blueprint document**: advice/full documentation to resource providers/users on how to engage with the federated virtualised environment. A living document on the EGI Wiki.
- **Test bed**: implement interfaces and services for a federated IaaS cloud on the basis of the Task Force blueprint and the available standards and technologies.

Task Force – Sept 2013

Cyfronet | FZJ | OeRC | EGI.eu | CESNET | GWDG | INFN-BARI

CNRS | IN2P3

KTH | Masaryk

FCTSG | INFN-CNAF

CETA | CESGA

IGI | SARA

RADICAL | IFCA

STFC | SZTAKI

BSC | GRNET

Imperial | DANTE

LMU | IPHC | IISAS | SixSq | 100%IT | CSC | IFAE | SRCE

Members

- ~70 individuals
- ~35 institutions
- >13 countries

Stakeholders

- 23 Resource Providers
  - 10 production
- 10 Technology Providers
- 8 User Communities
- 4 Liaisons

Technologies

- OpenNebula.
- StratusLab.
- OpenStack.
- Synnefo.
- WNoDeS.
- PERUN
- SlipStream

EGI's Core Infrastructure

The services that federate and integrate the functional services deployed in the **production infrastructure**

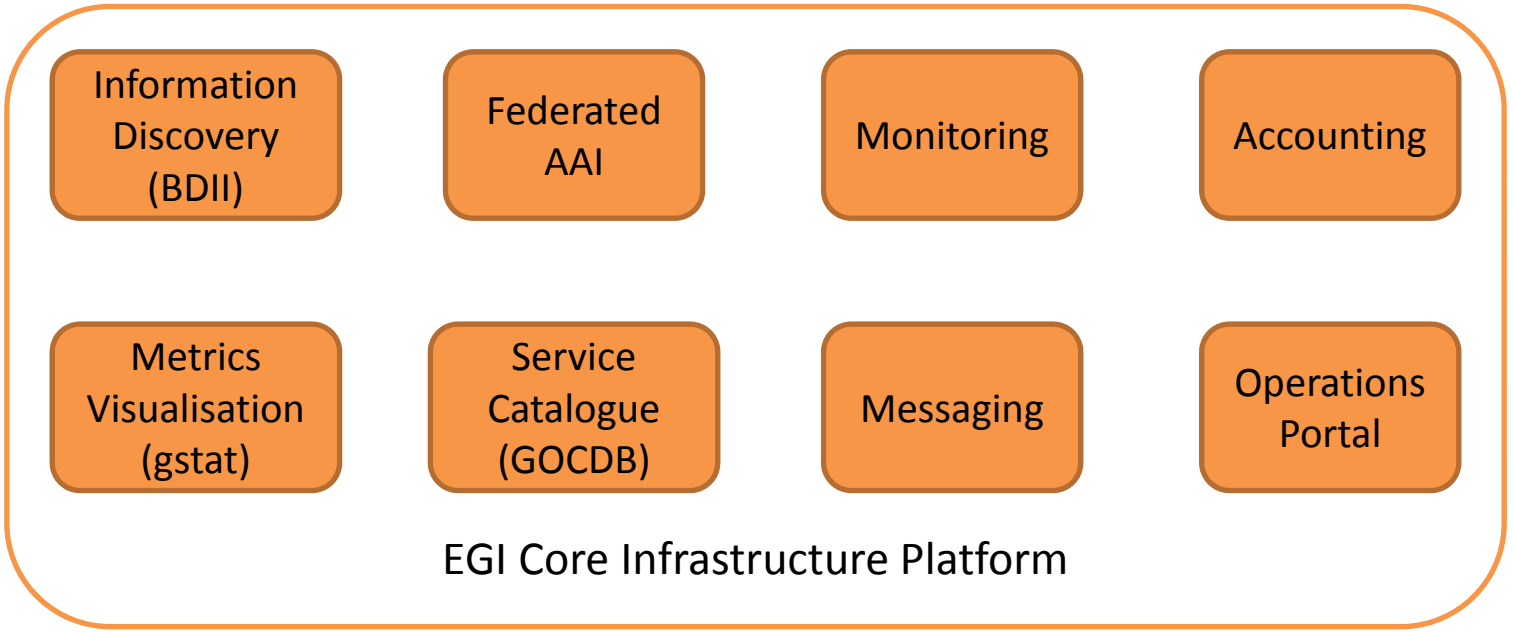

Services

Information system. Each cloud infrastructure exposes a LDAP server publishing information by means of an extended OGF GLUE2 schema. Each LDAP server is polled by a top-BDII server.

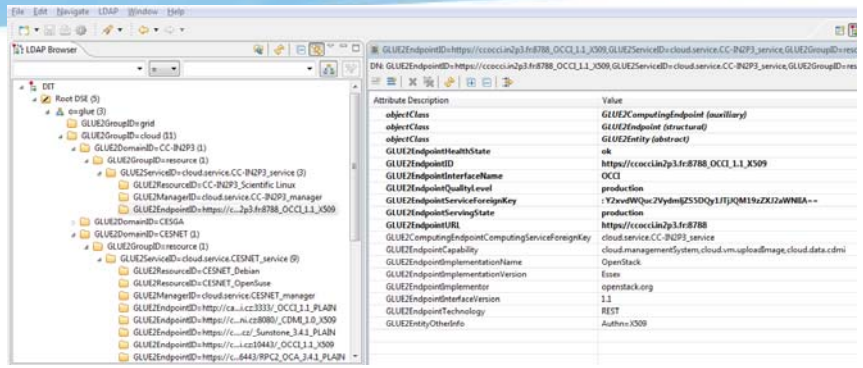

Monitoring. A standard Nagios installation is used to monitor the availability of the management interfaces exposed by each cloud infrastructure. Probes to test the state of the federated services are under development. https://cloudmon.egi.eu/nagios/

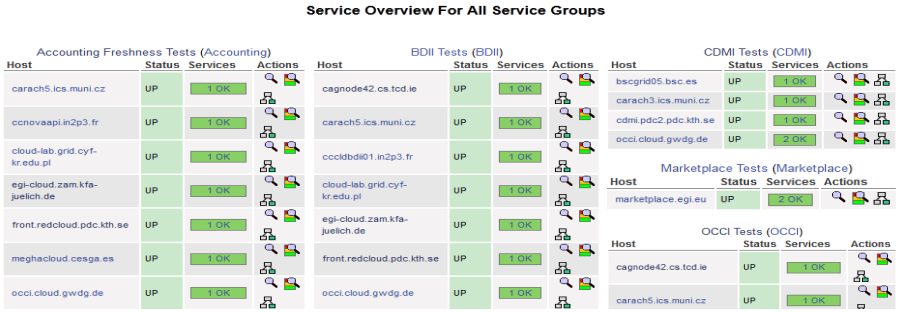

Service Overview For All Service Groups

Accounting Freshness Tests (Accounting)

Host	Status	Services	Actions
carach5.ics.muni.cz	UP	1 OK	
ccnovaapi.in2p3.fr	UP	1 OK	
cloud-lab.grid.cyf-kr.edu.pl	UP	1 OK	
egi-cloud.zam.kfa-juelich.de	UP	1 OK	
front.redcloud.pdc.kth.se	UP	1 OK	
meghacloud.cesga.es	UP	1 OK	
occi.cloud.gwdg.de	UP	1 OK	

BDII Tests (BDII)

Host	Status	Services	Actions
cagnode42.cs.tcd.ie	UP	1 OK	
carach5.ics.muni.cz	UP	1 OK	
cccldbdii01.in2p3.fr	UP	1 OK	
cloud-lab.grid.cyf-kr.edu.pl	UP	1 OK	
egi-cloud.zam.kfa-juelich.de	UP	1 OK	
front.redcloud.pdc.kth.se	UP	1 OK	
occi.cloud.gwdg.de	UP	1 OK	

CDMI Tests (CDMI)

Host	Status	Services	Actions
bscgrid05.bsc.es	UP	1 OK	
carach3.ics.muni.cz	UP	1 OK	
cdmi.pdc2.pdc.kth.se	UP	1 OK	
occi.cloud.gwdg.de	UP	2 OK	

Marketplace Tests (Marketplace)

Host	Status	Services	Actions
marketplace.egi.eu	UP	2 OK	

OCCI Tests (OCCI)

Host	Status	Services	Actions
cagnode42.cs.tcd.ie	UP	1 OK	
carach5.ics.muni.cz	UP	1 OK	

Accounting. Each cloud infrastructure generates usage records based on an extended version of the OGF UR format recommendation. Records are uploaded to a central server by means of a client customised for each type of infrastructure.

List of records contained in the cloud accounting database (last day).

Page last updated: 2012-09-18 22:00:03.971091

RecordId	Site	ZoneName	MachineName	Status	StartTime	EndTime	Network in (GB)	Network out (GB)	Memory (MB)	Disk (GB)	ImageId	
2012-09-17 21:00:01+00:00 CESNET vm-0	CESNET	EU	'one-0'	completed	2011-10-17 08:31:04	2011-10-17 10:41:16	0	2	512	None	None	C
2012-09-17 21:00:01+00:00 CESNET vm-1	CESNET	EU	'one-1'	completed	2011-10-17 10:46:45	2011-10-17 11:10:17	0	0	512	None	None	C
2012-09-17 21:00:01+00:00 CESNET vm-10	CESNET	EU	'hmmm_3'	completed	2011-10-18 12:45:15	2011-10-18 13:58:41	0	14	512	None	None	C
2012-09-17 21:00:01+00:00 CESNET vm-10440	CESNET	EU	'one-10440'	completed	2012-06-23 16:19:54	2012-06-23 16:26:08	0	0	256	None	None	C

http://accounting-devel.egi.eu/cloud.php
http://goc-accounting.grid-support.ac.uk/cloudtest/cloudsites2.html

Network implications for services

- Minimal...
- High reliability
- Multiple providers (some NREN connected others currently not)

- May make use of other services provided e.g. Fed ID though many competitors in this area

EGI's Cloud Infrastructure

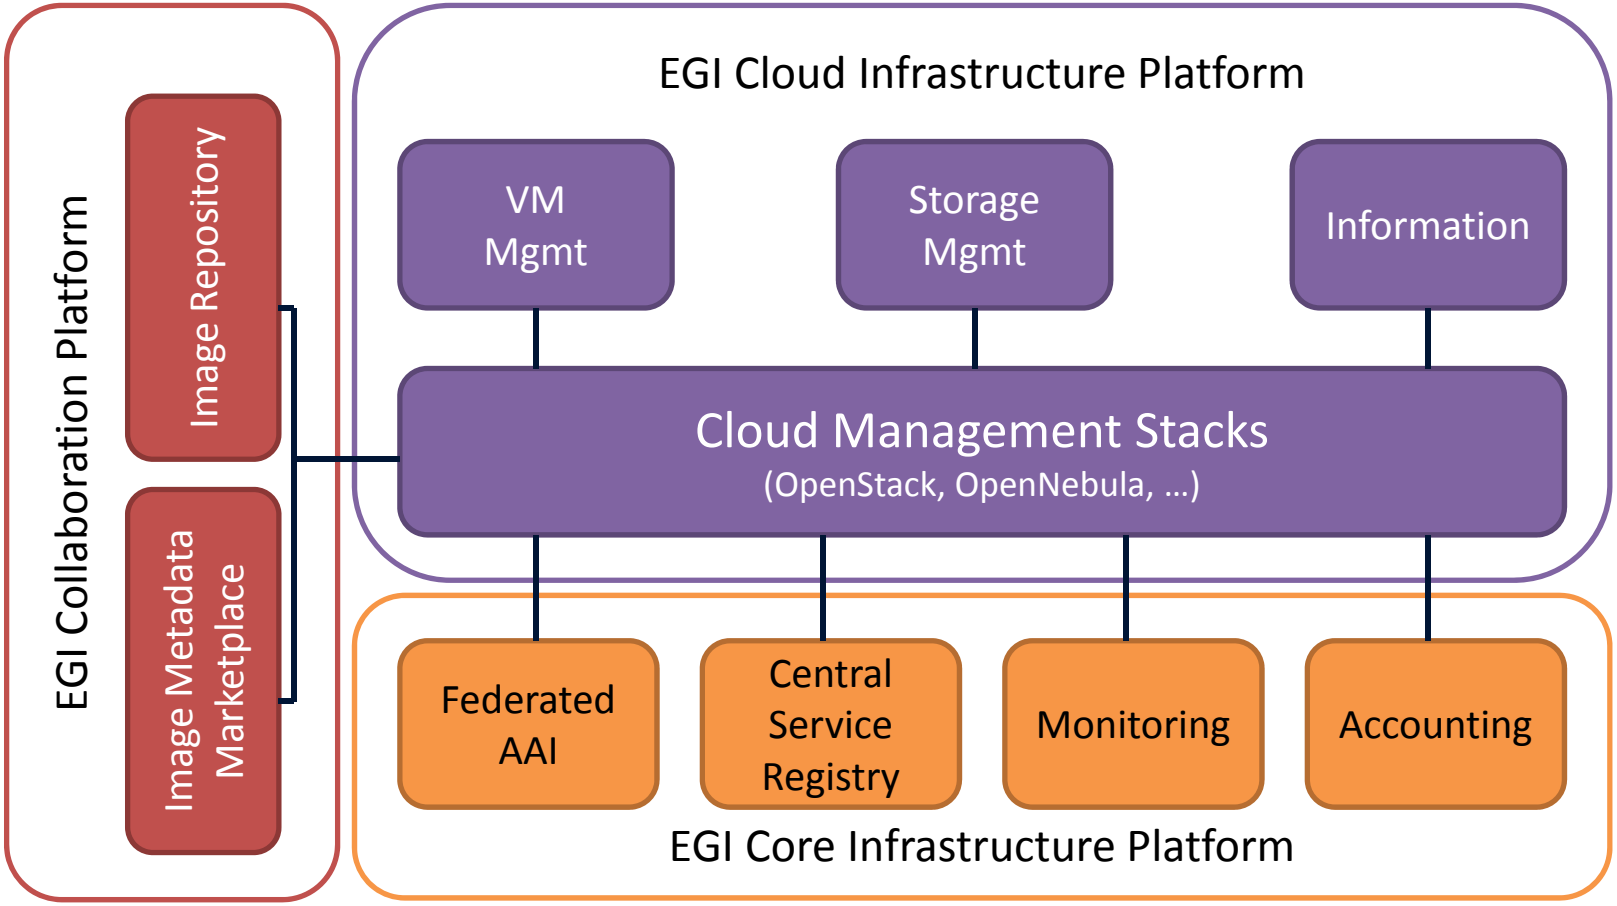

Cloud Services

Marketplace. A repository were Resource Providers and EGI can publish metadata about images from which virtual machines can be instantiated. When needed, a single image can be signed and then endorsed by multiple providers.

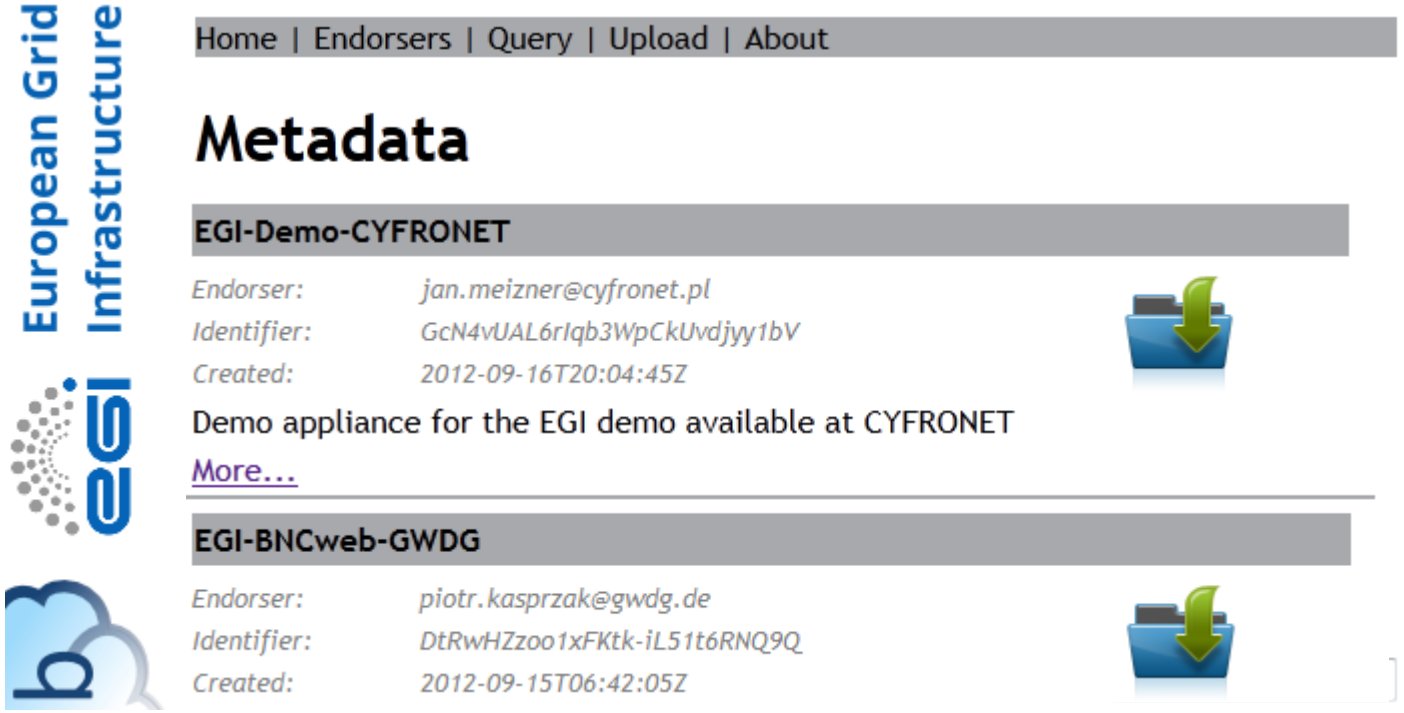

http://marketplace.egi.eu/metadata/

FedCloud Usage Models

NB: Users = researchers

IaaS for power users

- Comparable experience to ec2-api-tools, euca2ools.
- Command line interface for easy scripting and integration with workflow.
- Multiple language bindings to create a developer-oriented ecosystem.

IaaS as a backend

- Promote uniform OCCI server implementations.
- Facilitate the development of connectors and shims.
- Platforms

IaaS for end users

- Portals. Upcoming IGI portal + SlipStream EGI mini project.
- Community services.

Demo Use Cases – September 2013

BNCweb (CLARIN)

- Interface to the British National Corpus (100ml words).
- Used in a wide variety of computational linguistic applications.
- Searching the text and exploiting the detailed textual metadata.
- Open source project, freely available for educational and research purposes.

Use cases

1. Teaching a masters course in 'Exploring English Usage' in October-November 2012, and 'Corpus Linguistics' in February-March 2013.
2. Federated search in the CLARIN European e-Infrastructure.
3. Build additional web services on top of BNCWeb, e.g. adding improved visualizations of the search results.

EU Federated IaaS Cloud for CLARIN and BNCWeb

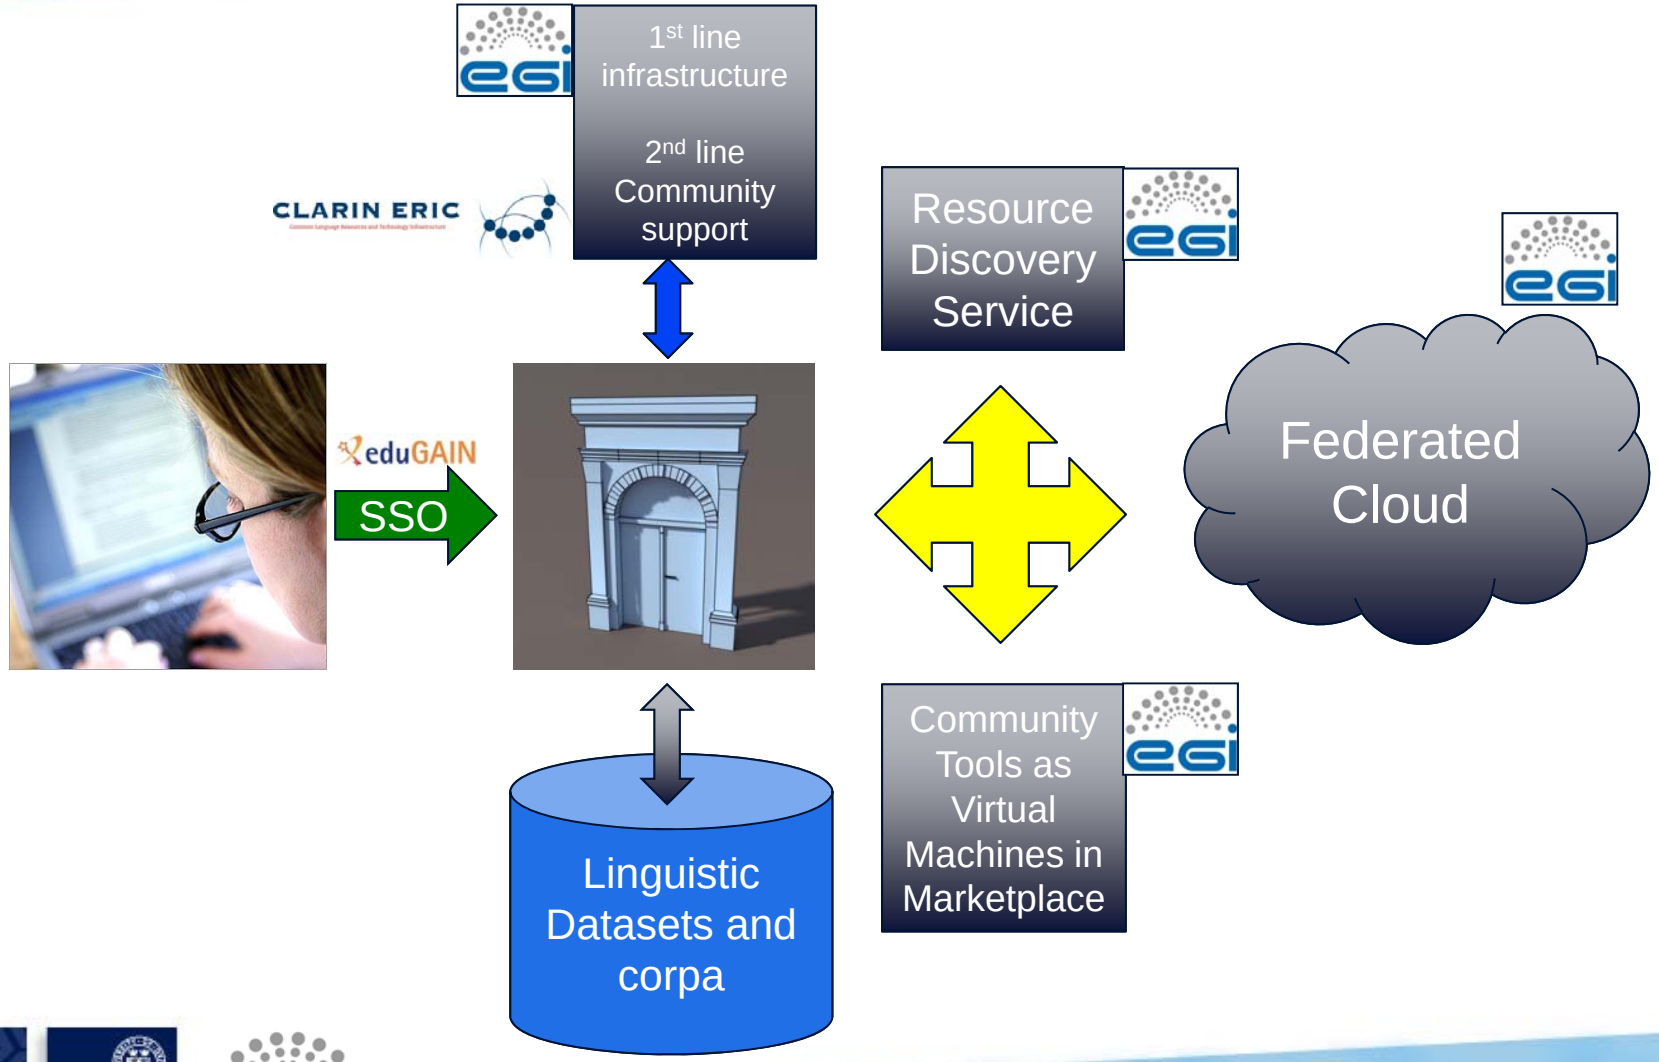

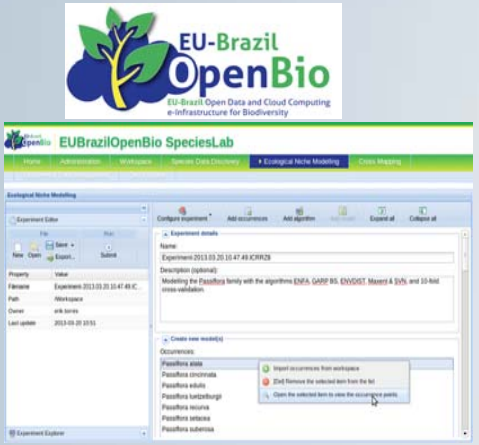

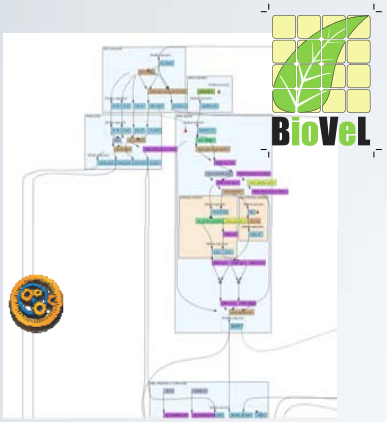

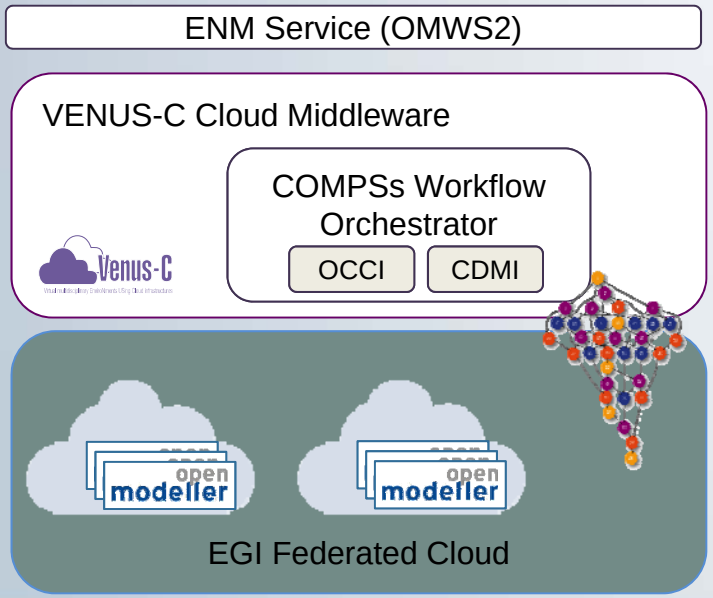

- **Shared requirements** between EUBrazilOpenBio and BioVeL
- The **EUBrazilOpenBio ENM service** is exposed through an extended openModeller Web Service interface (**OMWS2** in the picture).
- Such interface in EUBrazilOpenBio supports multi-staging and multiparametric experiments implemented through **COMPSs** and the openModeller software and managed through a Virtual Research Environment (VRE) portal.
- The OMWS extensions are backwards compatible with the original specification, allowing existing clients, as the Taverna Workflow Management System in BioVeL, to be fully supported in the new implementation.
- In the case of the EGI Federated Cloud, the **VENUS-C** middleware is used to instantiate openModeller workflows on **cloud resources** from different providers **dynamically deployed by COMPSs**.
- An **OCCI connector** is used for the VMs management while data management supports **CDMI** endpoints.

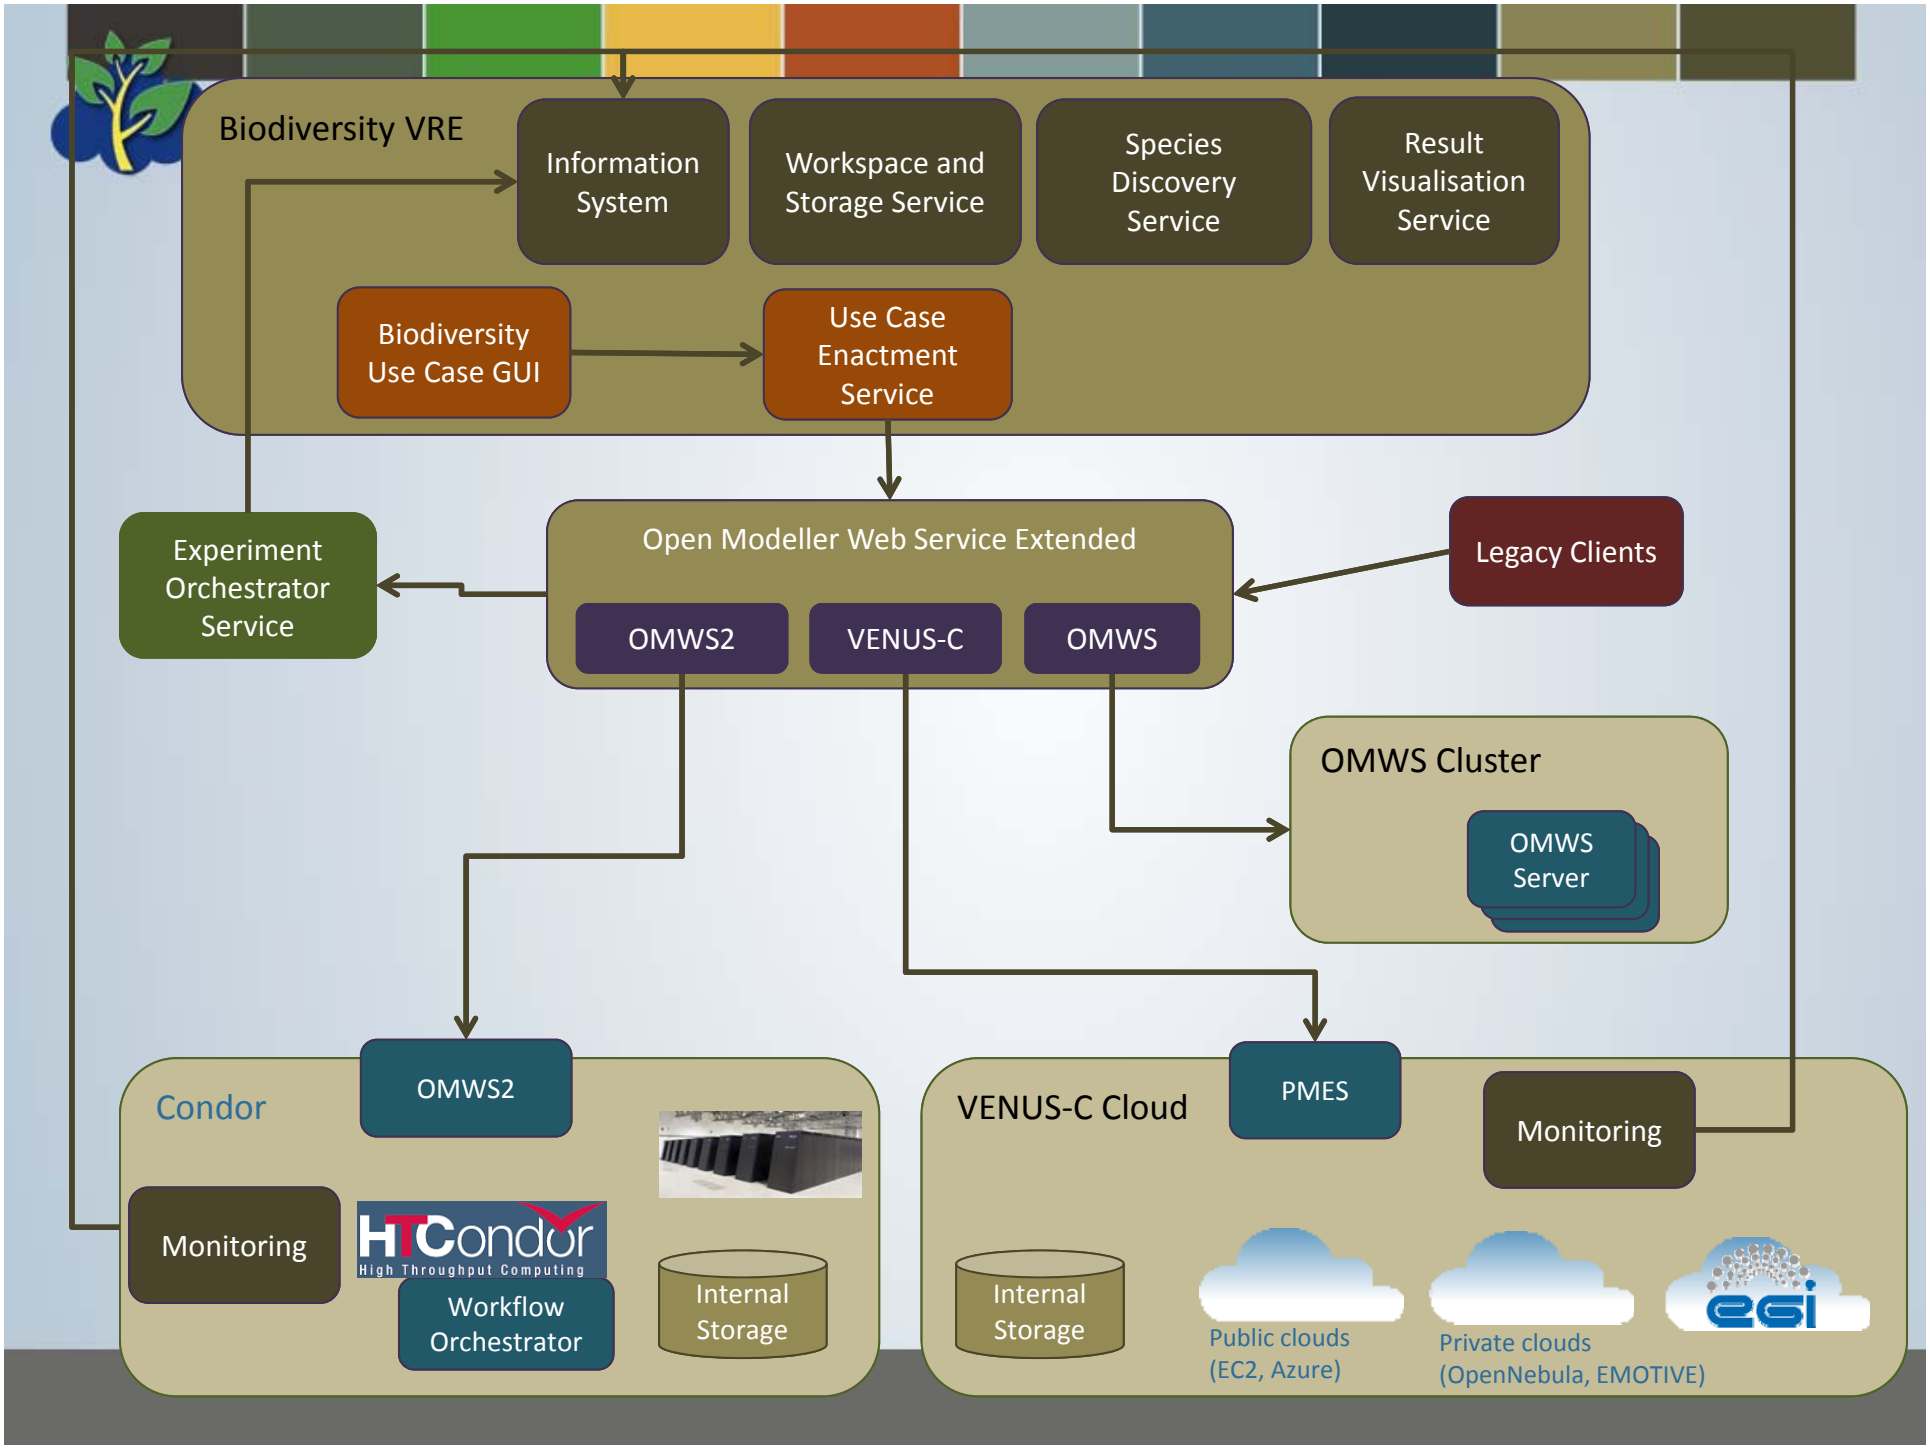
Biodiversity VRE
Information System
Workspace and Storage Service
Species Discovery Service
Result Visualisation Service
Biodiversity Use Case GUI
Use Case Enactment Service
Experiment Orchestrator Service
Open Modeller Web Service Extended
OMWS2
VENUS-C
OMWS
Legacy Clients
OMWS Cluster
OMWS Server
Condor
OMWS2
Monitoring
HTCondor
High Throughput Computing
Workflow Orchestrator
Internal Storage
VENUS-C Cloud
PMES
Monitoring
Internal Storage
Public clouds (EC2, Azure)
Private clouds (OpenNebula, EMOTIVE)
egi

Use Cases

WeNMAR – Structural Biology

- Validating biomolecular NMR structures. VM images pre-packaged with VirtualCing and ToPoS.
- Training. VM images pre-packaged with Gromacs.

ESA

- Portal interface, with custom cluster deployment
- Instantiate an ESA configured Computing Cluster on the Federated Cloud
- Start processing jobs on the cluster

WS-PGRADE/gUSE – Software Engineering

- User support environment for grid and clouds + GUI.
- Flexible deployment of virtualised DCI to run predefined, repeatable tests.

Network implications for usecases

- Not all resources are public sector or academic, the # of these will only grow
- Researchers expect performance to be uniform across all resources within the cloud
- Large scale public cloud -> normally a single migration of data into a single cloud
- Federated cloud -> multiple clouds, multiple data movements
- Several usecases involve low latency requirements due to service compositions
- In future on demand dynamic network configuration would be advantageous for usecases involving sensitive data
  - We would want these to be achieved using open standards where possible rather than proprietary services

EGI's Grand Vision

Establish by 2020 a distributed open compute and data infrastructure comprising a 10M Core Federated Cloud and 10 Exabyte of Federated Cloud Storage across Europe that is able to support the data analysis activities of all researchers within the European Research Area

EGI Council Discussion – June 2013

Conclusions

- Working towards production-grade cloud resources within the EGI infrastructure
- Currently working with resource providers to iron out small network limitations in last mile
- Biggest challenge…

Data

- Focusing on use cases and early adopters:
  - dedicated user support;
  - policies for tailored image management and support;
  - SLA.

Demo Team

Viet Tran

Matteo Turilli

Zeeshan Ali Shah

Sandor Acs

Alvaro Lopez

Stuart Pullinger

Björn Hagemeier

Emir Imamagic

Alison Packer

Piotr Kasprzak

Owen Synge

Mattieu Puel

Ivan Diaz Alvarez

Michel Drescher

Mohammed Airaj

Salvatore Pinto

Jan Meizner

Peter Solagna

Boris Parak

Daniele Lezzi

Alvaro Simon

John Gordon

Kostas Koumantaros

Elisabetta Ronchieri

Andy Edmonds

Florian Feldhaus

Ilja Livenson

Roberto Rosende Dopazo

Michel Jouvin


Task Force resources

- Mailing List: fedcloud-tf@mailman.egi.eu
- Wiki site: http://go.egi.eu/tf-fedclouds
- GitHub: https://github.com/EGI-FCTF
- Indico site: https://www.egi.eu/indico/categoryDisplay.py?categId=56
- Contact the chair: david.wallom@oerc.ox.ac.uk

Thank you