One year of The EGI Federated Clouds Task Force

EGI Technical Forum 2012 – Prague, 17/21 Sep

Matteo Turilli

Senior Research Associate, OeRC, University of Oxford
Chair – EGI Federated Clouds Task Force
matteo.turilli@oerc.ox.ac.uk

Outline

- EGI new challenges and cloud computing.
- TF objectives, deliverables, mandate and membership.
- Federation model.
- Federation test bed and test bed demos.
- Blueprint document, joining procedure and know how.
- From Task Force to Task within EGI-InSPIRE.
- Use cases.
- Conclusions.

EGI New Challenges and Cloud Computing

- 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.

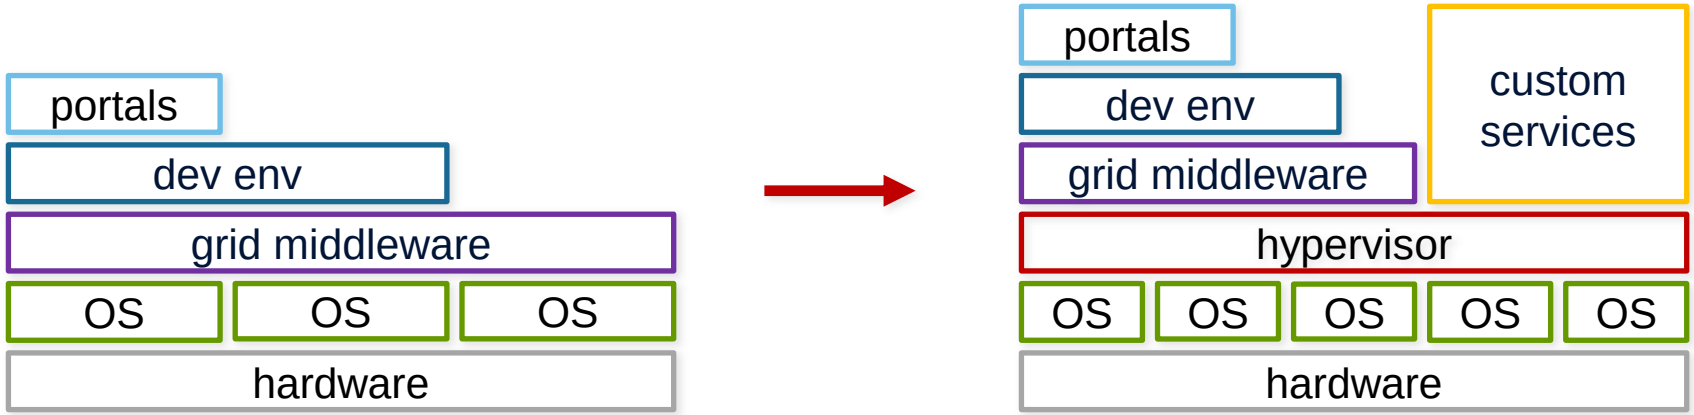

EGI New Challenges and Cloud Computing

Personalised environments for individual research communities in the European Research Area.

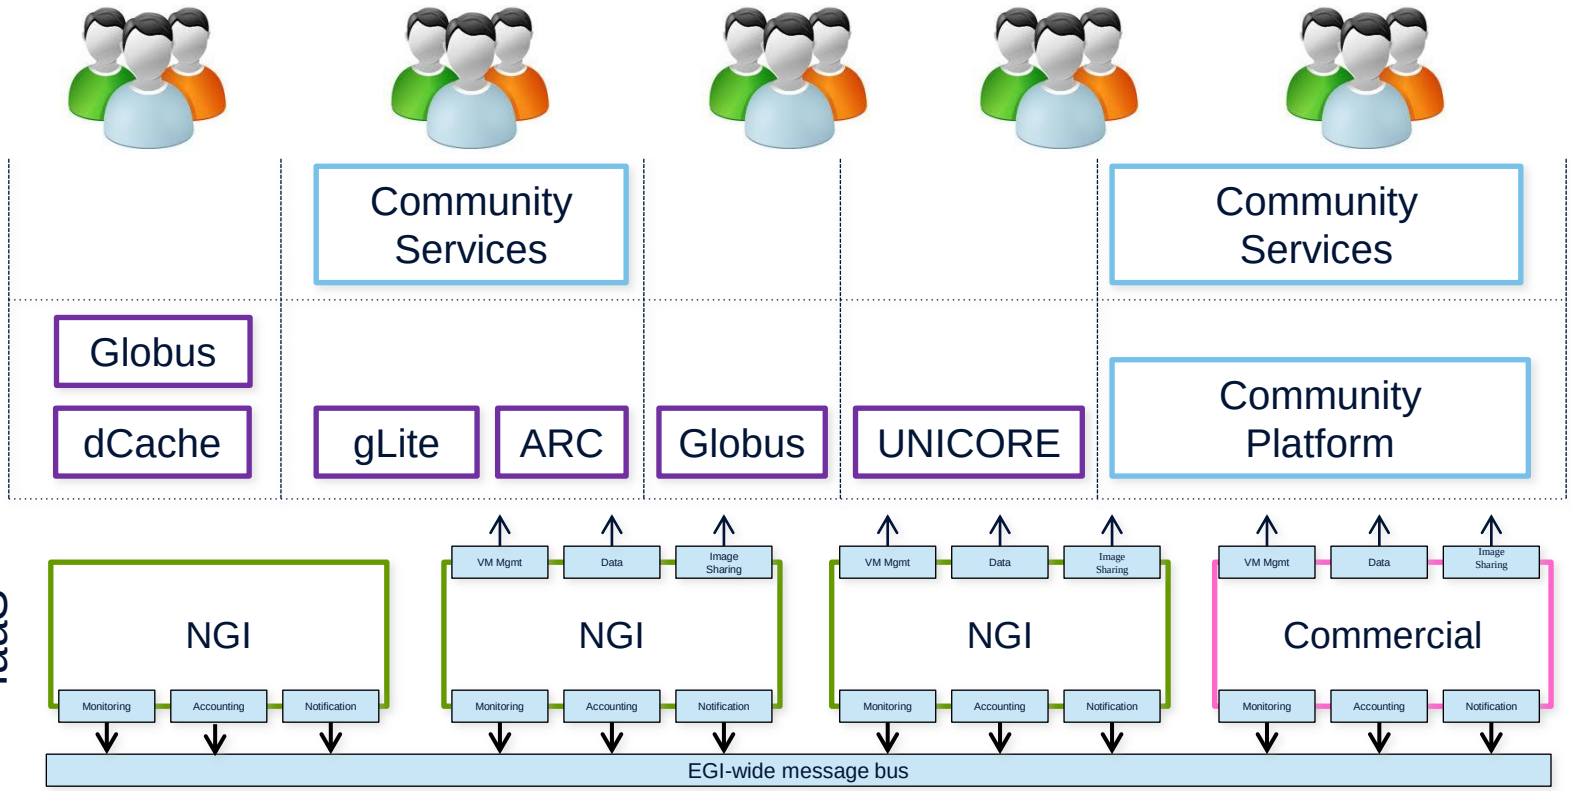

TF Objectives and Deliverables

- **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 cloud on the basis of the Task Force blueprint and the available standards and technologies.

Task Force Mandate and Organisation

Mandate: 18 months, September 2011 – March 2013.

Activities: 3 blocks of 6 months each.

1. **Setup**: Sep 2011 – Mar 2012.
   - Engagement of resource and technology providers.
   - Federation model.
   - Draft of the Blueprint document and demo.
2. **Consolidation**: Mar 2012 – Sep 2012.
   - Engagement of user communities.
   - Test bed and first use case.
   - Draft of the Blueprint document and demo.
3. **Integration**: Sep 2012 – Mar 2013.
   - Test bed and early adopters.
   - Publication of the Blueprint document and demo.

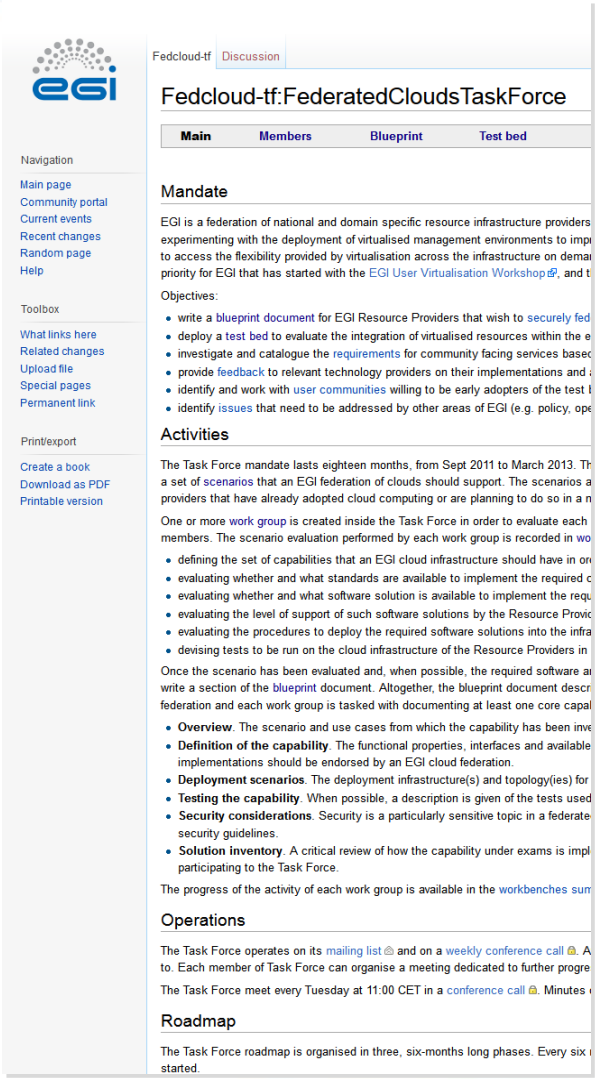

Fedcloud-tf | Discussion

Fedcloud-tf:FederatedCloudsTaskForce

Main | Members | Blueprint | Test bed

Navigation
Main page
Community portal
Current events
Recent changes
Random page
Help

Toolbox
What links here
Related changes
Upload file
Special pages
Permanent link

Print/export
Create a book
Download as PDF
Printable version

Mandate

EGI is a federation of national and domain specific resource infrastructure providers
experimenting with the deployment of virtualised management environments to impr
to access the flexibility provided by virtualisation across the infrastructure on dema
priority for EGI that has started with the EGI User Virtualisation Workshop, and t

Objectives:

- write a blueprint document for EGI Resource Providers that wish to securely fed
- deploy a test bed to evaluate the integration of virtualised resources within the e
- investigate and catalogue the requirements for community facing services basec
- provide feedback to relevant technology providers on their implementations and
- identify and work with user communities willing to be early adopters of the test
- identify issues that need to be addressed by other areas of EGI (e.g. policy, ope

Activities

The Task Force mandate lasts eighteen months, from Sept 2011 to March 2013. Th
a set of scenarios that an EGI federation of clouds should support. The scenarios a
providers that have already adopted cloud computing or are planning to do so in a n

One or more work group is created inside the Task Force in order to evaluate each
members. The scenario evaluation performed by each work group is recorded in wo

- defining the set of capabilities that an EGI cloud infrastructure should have in or
- evaluating whether and what standards are available to implement the required c
- evaluating whether and what software solution is available to implement the requ
- evaluating the level of support of such software solutions by the Resource Provic
- evaluating the procedures to deploy the required software solutions into the infra
- devising tests to be run on the cloud infrastructure of the Resource Providers in

Once the scenario has been evaluated and, when possible, the required software ar
write a section of the blueprint document. Altogether, the blueprint document descr
federation and each work group is tasked with documenting at least one core capal

- **Overview**. The scenario and use cases from which the capability has been inve
- **Definition of the capability**. The functional properties, interfaces and available
  implementations should be endorsed by an EGI cloud federation.
- **Deployment scenarios**. The deployment infrastructure(s) and topology(ies) for
- **Testing the capability**. When possible, a description is given of the tests used
- **Security considerations**. Security is a particularly sensitive topic in a federate
  security guidelines.
- **Solution inventory**. A critical review of how the capability under exams is impl
  participating to the Task Force.

The progress of the activity of each work group is available in the workbenches sum

Operations

The Task Force operates on its mailing list and on a weekly conference call. A
to. Each member of Task Force can organise a meeting dedicated to further progre

The Task Force meet every Tuesday at 11:00 CET in a conference call. Minutes

Roadmap

The Task Force roadmap is organised in three, six-months long phases. Every six
started.

Task Force Members and Technologies

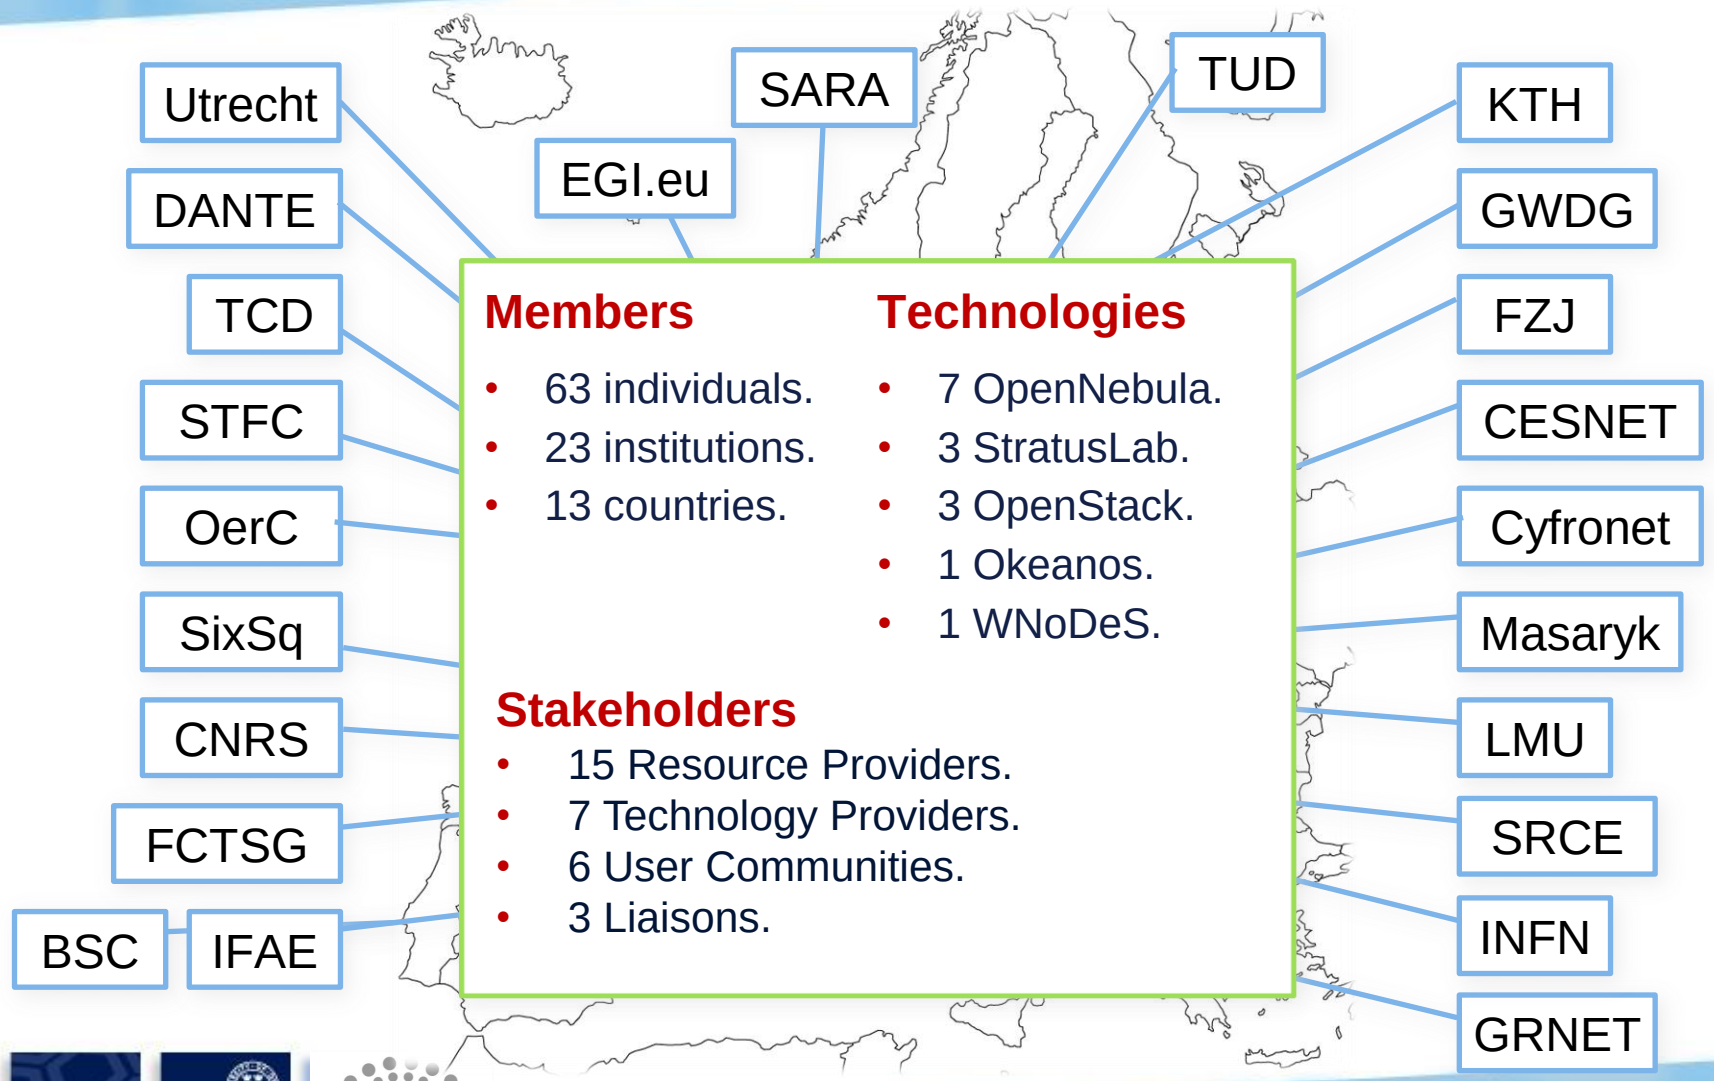

Members

- 63 individuals.
- 23 institutions.
- 13 countries.

Technologies

- 7 OpenNebula.
- 3 StratusLab.
- 3 OpenStack.
- 1 Okeanos.
- 1 WNoDeS.

Stakeholders

- 15 Resource Providers.
- 7 Technology Providers.
- 6 User Communities.
- 3 Liaisons.

Federation Model

Cloud initiatives and landscape ~Sep 2011

- Early development stage of open source solutions.
- Prevalence of test bed and/or pilot projects for cloud evaluation by resource providers.
- Multiple cloud management platform with proprietary interfaces.
- At best, early stage of user requirements elicitation.
- On-going security policy evaluation, early stage of integration with IT departments.

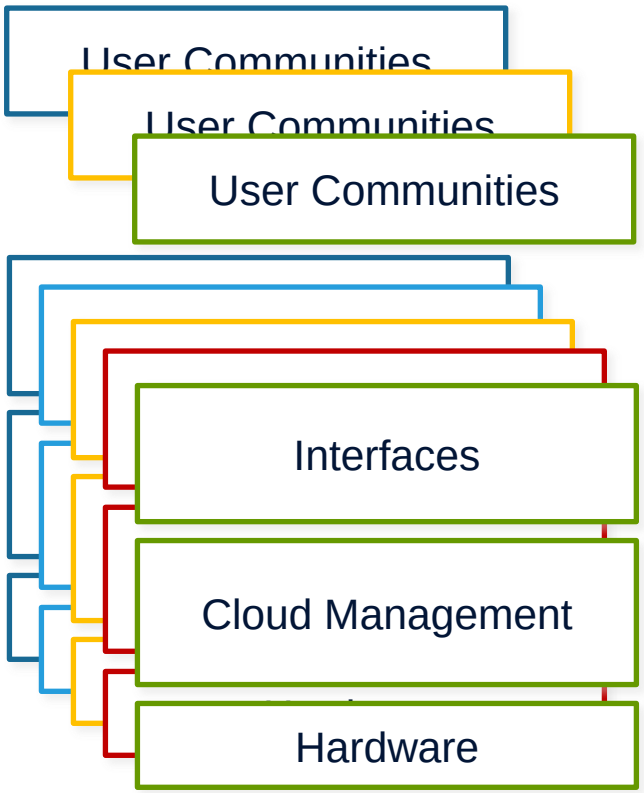

Federation Model

"EGI Cloud Integration Profile", S. Newhouse, M. Drescher.

- **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 cloud technology.
- **Provider agnosticism**: the only condition to federate resources is to expose the chosen interfaces and services.

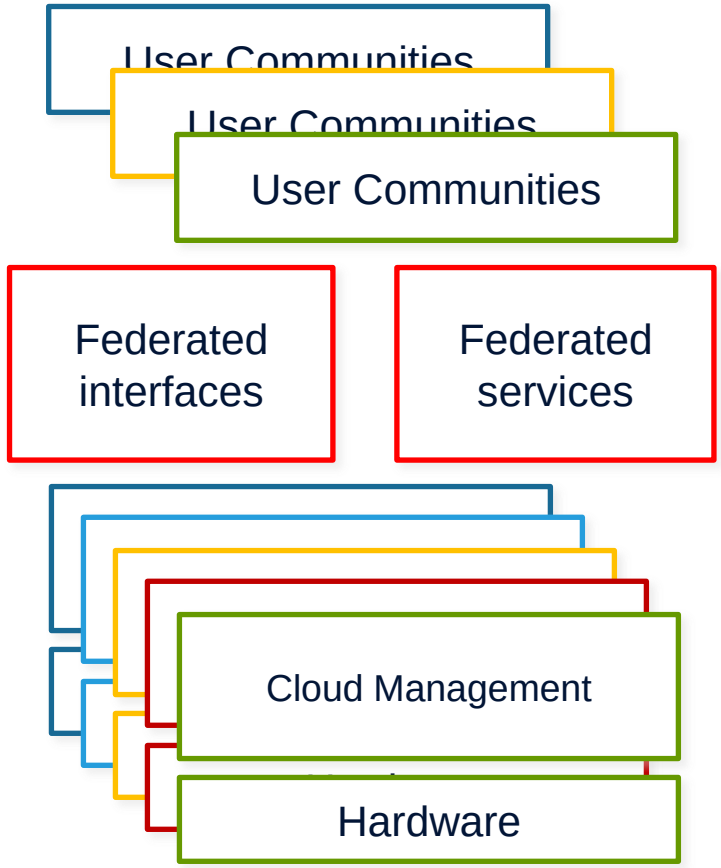

Workgroups and Workbenches

Interfaces

VM Management
OCCI 1.1

Data Management
CDMI

Services

Information Systems
GLUE2+, LDAP/BDII

Monitoring
Nagios, cloud probes

VMs repository
Marketplace

Accounting
UR+, Apel

Federated AAI
X509, VOs

Clients and Brokering
rOCCI, CompatibleOne

Federation Test bed – Sep 2012

Composed of 4 services, 2 management interfaces, 7 cloud infrastructures operated by 6 Resource Providers. 3 more providers are in the process of being federated.

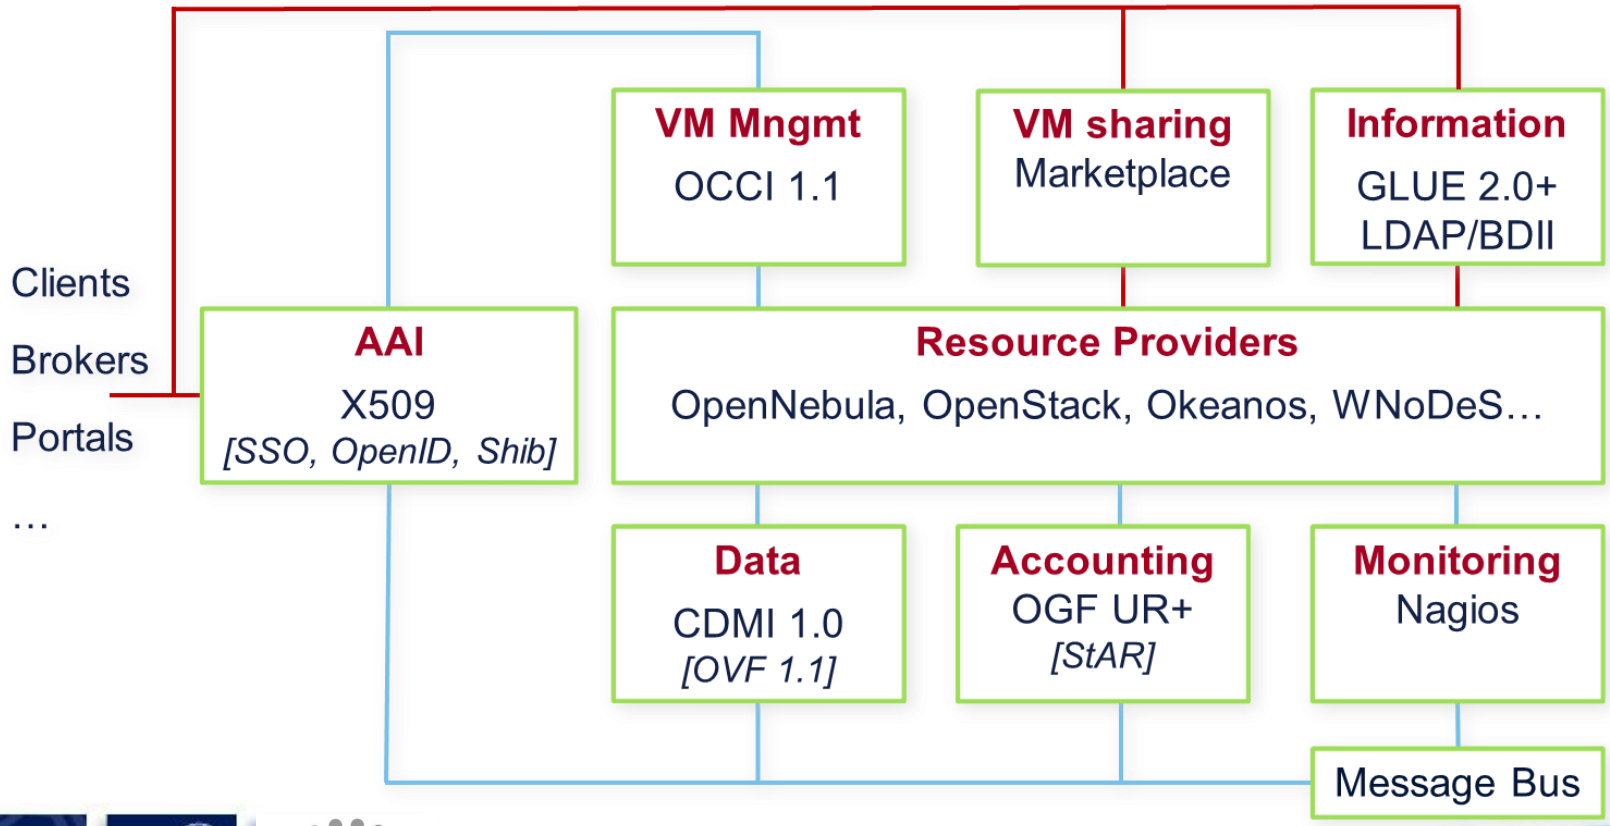

Services

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.

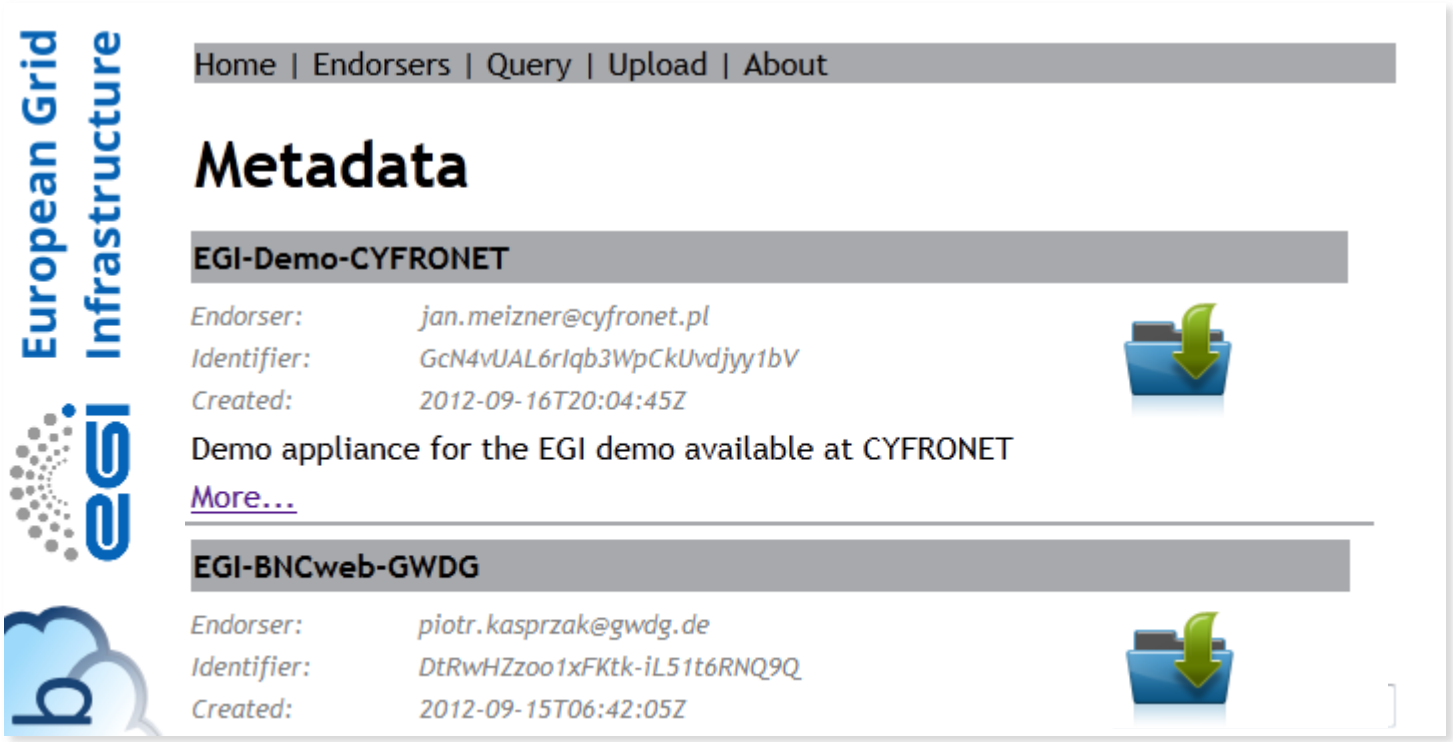

http://marketplace.egi.eu/metadata/

Services

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

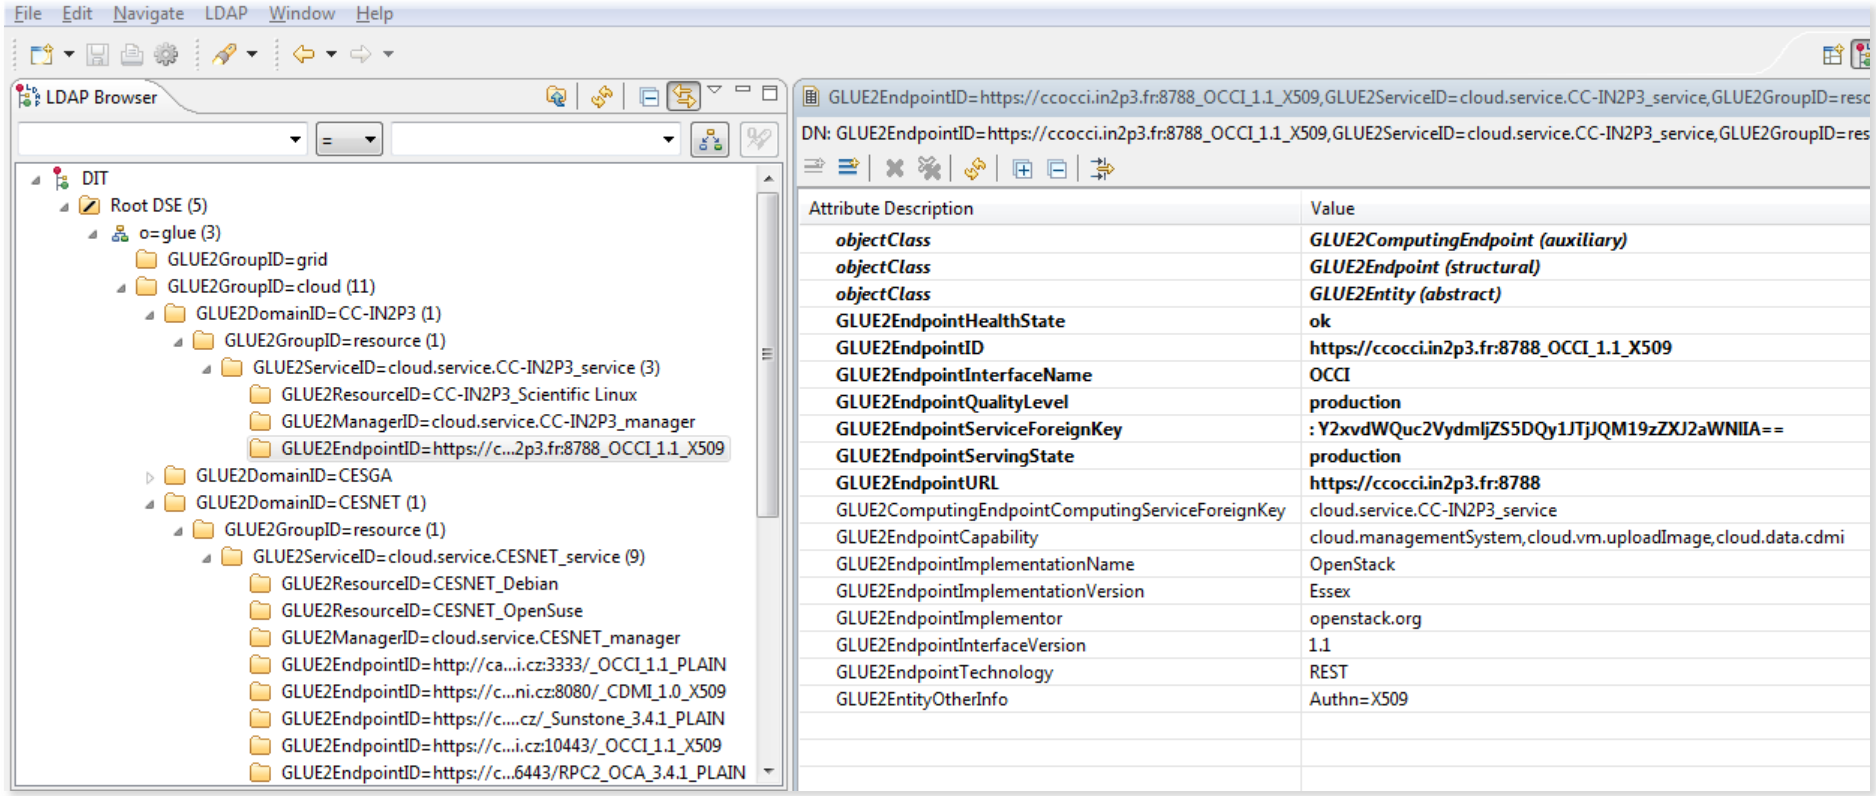

ldap://test03.egi.cesga.es:2170

Services

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.

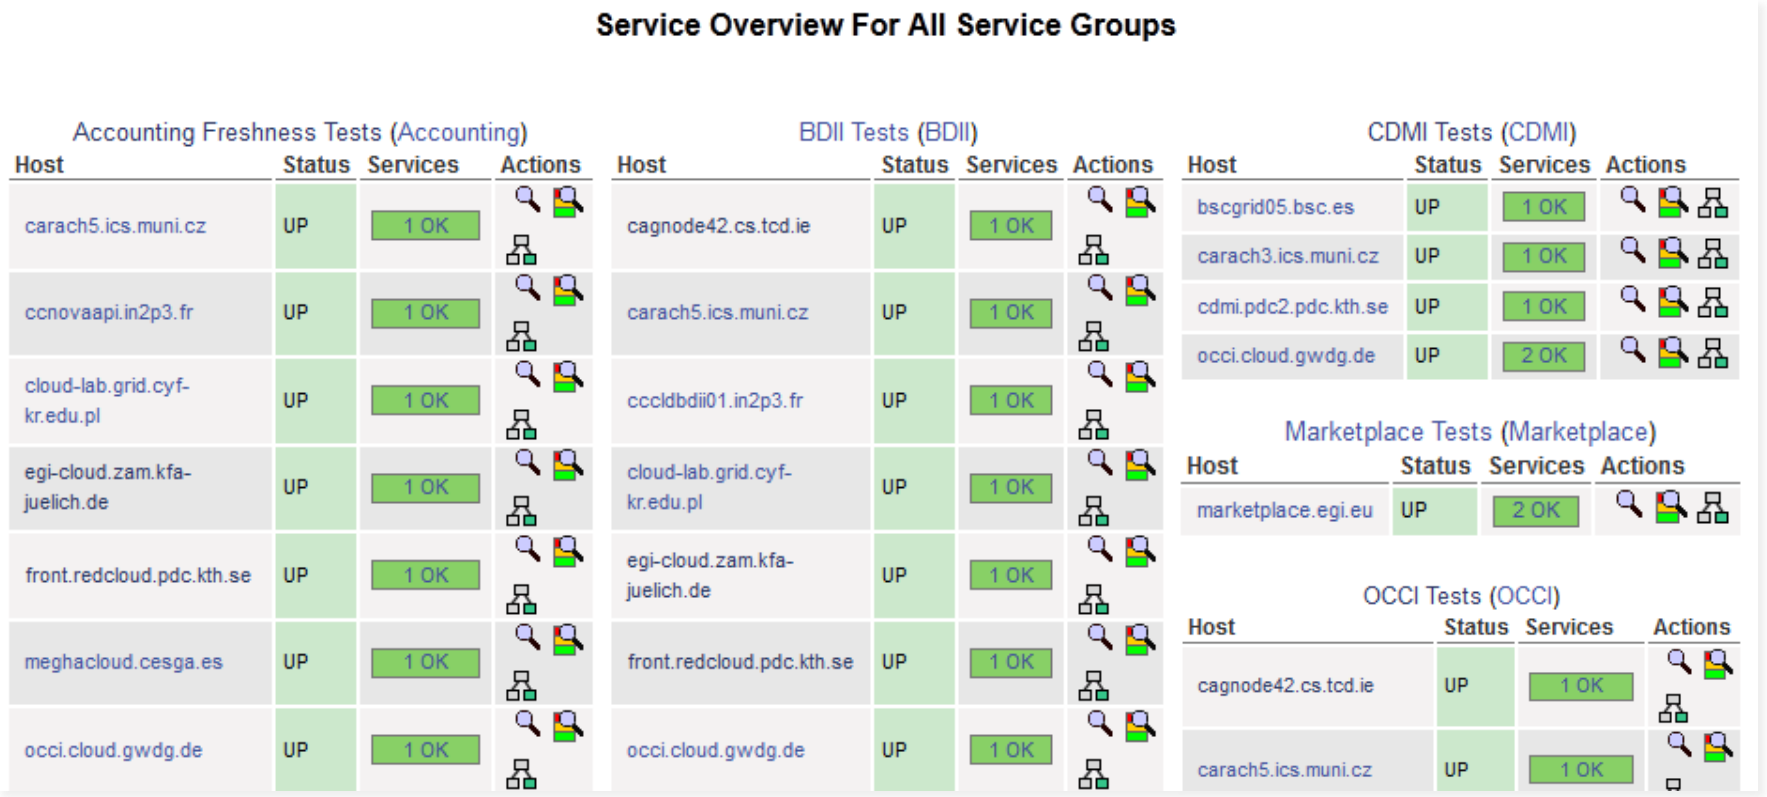

Service Overview For All Service Groups

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	

https://test30.egi.cesga.es/nagios/

Services

Accounting. Each cloud infrastructure generates usage records based on an extended version of the EGI 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	CloudType
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	OpenNebula
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	OpenNebula
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	OpenNebula
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	OpenNebula

http://goc-accounting.grid-support.ac.uk/cloudtest/vms.html

Federation Demo – Mar 2012

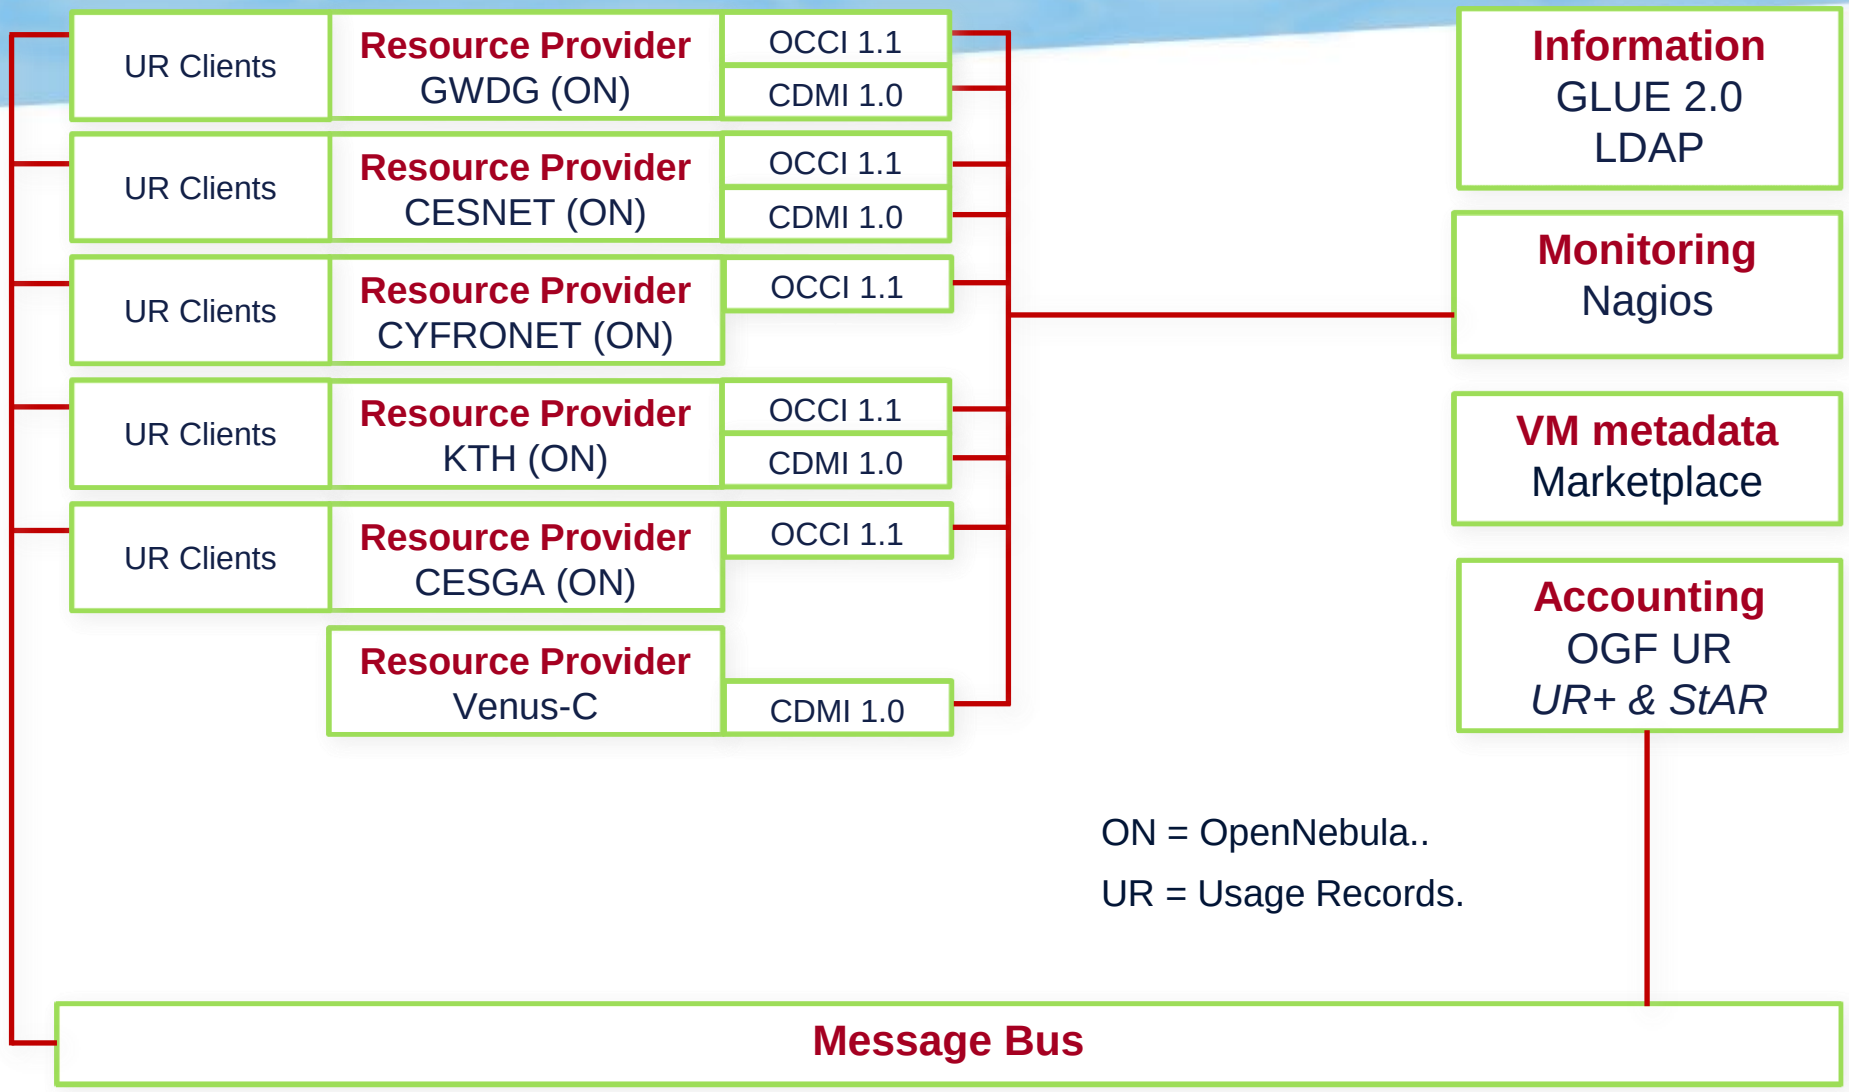

Federation Demo – Sep 2012

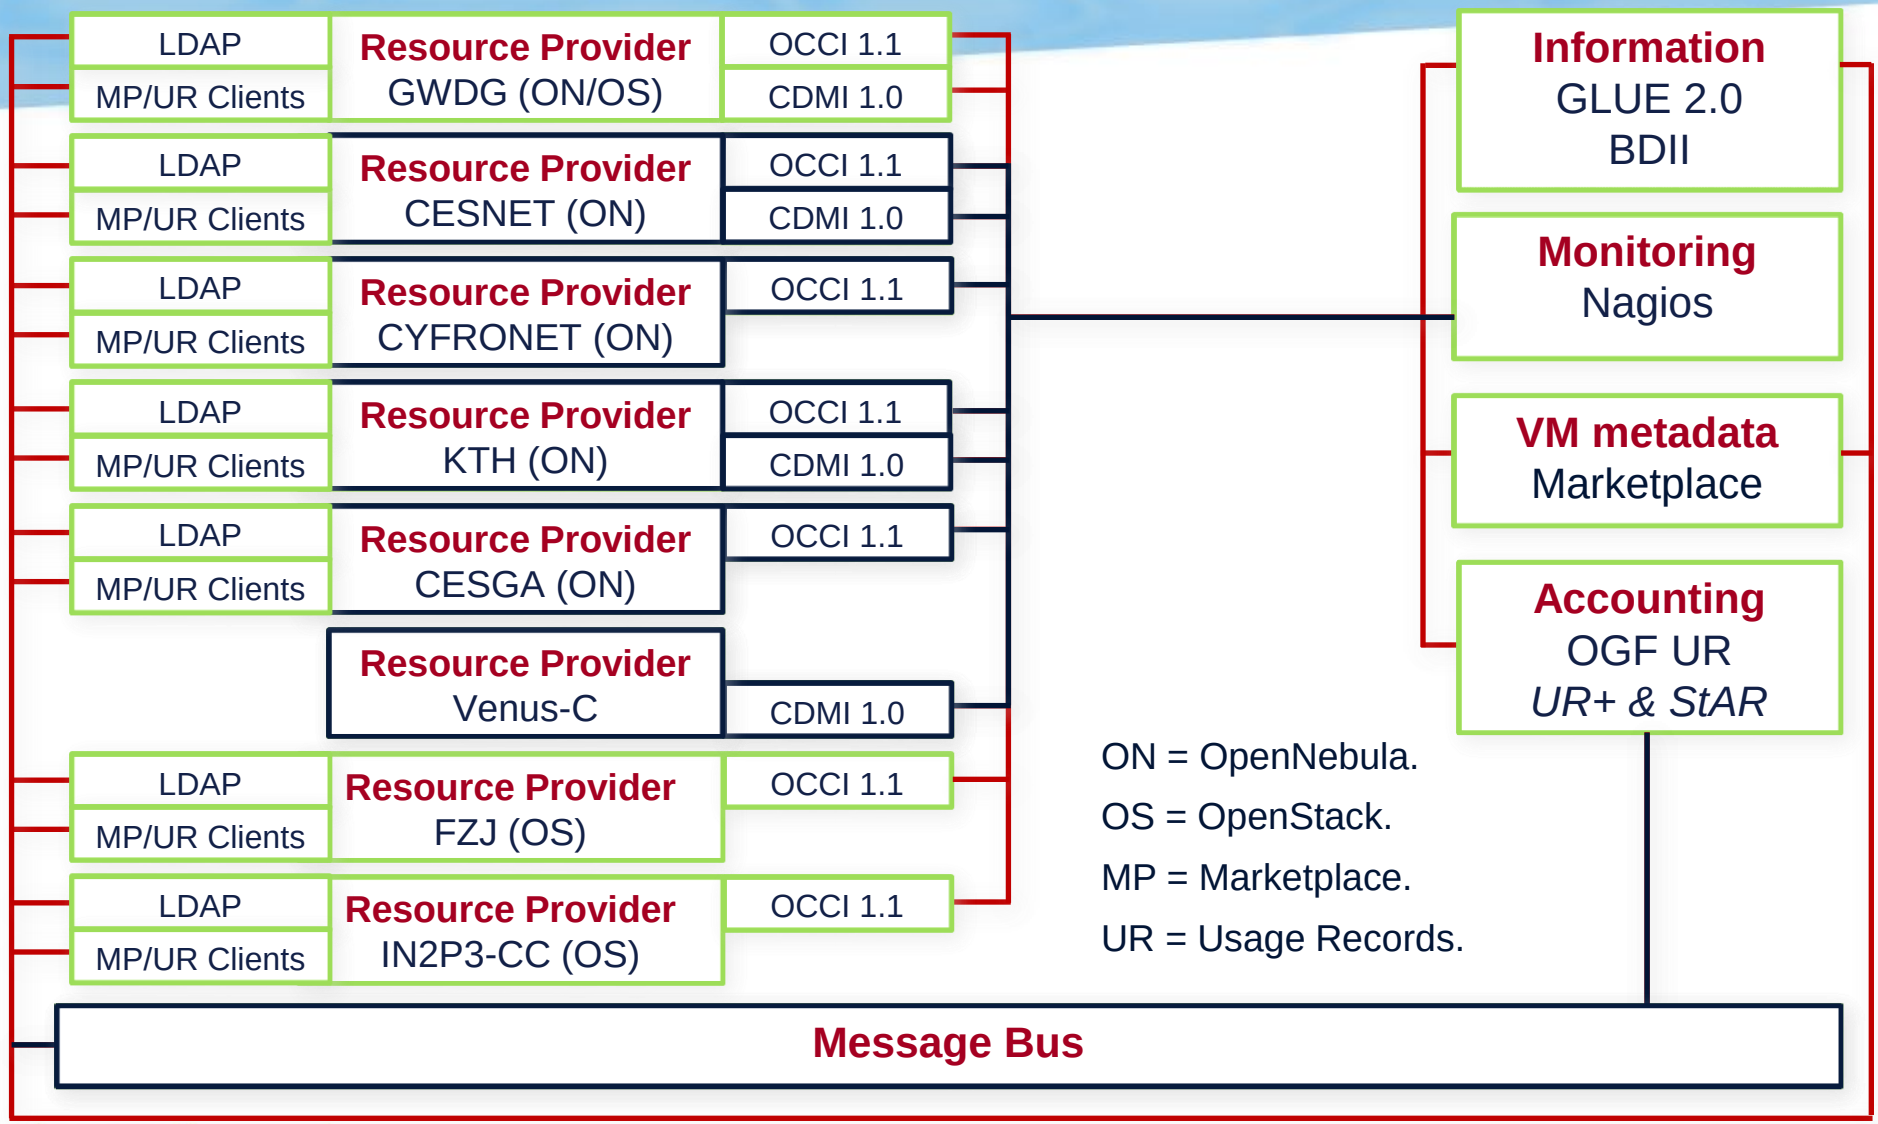

Blueprint Document

Fedcloud-tf | Discussion | Read | Edit | View history | Go | Search

Fedcloud-tf:Blueprint:Introduction

(Redirected from Fedcloud-tf:Blueprint)

Navigation
Main page
Community portal
Current events
Recent changes
Random page
Help

Toolbox
What links here
Related changes
Upload file
Special pages
Permanent link

Print/export
Create a book
Download as PDF

Main | Members | Blueprint | Test bed | Work groups | User Communities | Outreach | Administrative

Introduction | Overview | Conclusions | References

Capabilities: VM management | Data management | Information discovery | Accounting | Monitoring | Notification | Authentication and Authorisation | VM sharing |

Introduction [edit]

This blueprint document will be assembled and written by the Task Force members providing EGI Resource Centres with all information necessary to provision computing and storage resources through a cloud management layer so that they can be easily and securely federated as part of the EGI production infrastructure.

High-level scenarios [edit]

The blueprint is based on a clollection of six high-level Cloud usage scenarios that were collected and validated at a number of occasions in the EGI community. Those usage scenarios represent six archetypical Cloud scenarios that are common (or perceived to be common) in a distributed, hybrid Cloud landscape provided by EGI resource centres. Associated with these six scenarios are six key capabilities that, when implemented on a technical level, ensure interoperabilty across individual Cloud providers.

The following sub-sections illustrate the six usage-scenarios providing an overview of the technical

Contents [hide]

https://wiki.egi.eu/wiki/Fedcloud-tf:Blueprint

Blueprint Document

Join the federation:

- Expose an OCCI interface.
- Install an LDAP server with a GLUE2 schema tailored for cloud resources.
- Allows the Nagios probes to monitor the interfaces and services.
- Upload usage records to the EGI centralised repository.
- Publish the image metadata into the federation Marketplace
- Install, if needed, a CDMI server.

Documented Knowledge for:

- OpenStack, OpenNebula installation and configuration.
- OCCI and CDMI
- Marketplace.
- Nagios probes for cloud resources
- GLUE2 and UR for cloud resources.
- Latest developments in cloud brokering and clients.
- User communities leveraging cloud computing.

From Task Force to Task within EGI-InSPIRE

May 1st 2012: the Task Force becomes an official Task within EGI-InSPIRE.

- **Effort**: Effort officially contributed by: IIAS, KTH, Jülich, LUH, SARA, CISC, INFN, CNRS, CESNET and OeRC.
- **EGI Technical Outreach team**: Collaboration with the EGI Technical Outreach team to coordinate and promote the support of Virtual Research Communities that need or could benefit from Cloud Computing.
- **Community Engagement**: Collaboration with the EGI Community Engagement to organise the TF community activities.
  - 1st EGI Federated Clouds TF PlugFest, July 12th/13th, Amsterdam.
  - Set up of a track for requirement gathering dedicated to Cloud Computing.

Use Cases

- **Structural biology** – We-NMR project:
  Gromacs training environments.
- **Musicology** – Peachnote project:
  music score search engine and analysis platform.
- **Linguistics** – CLARIN project:
  scalable ‘British National Corpus’ service (BNCWeb).
- **Ecology** – BioVel project:
  remote hosting of OpenModeller service.
- **Software development** – SCI-BUS project:
  simulated environments for portal testing.
- **Space science** – ASTRA-GAIA project:
  data integration with scalable workflows.

Conclusions

Output

- Adoptions of standards for VM and data management.
- Federation model compatible and consistent with current EGI infrastructure.
- Contribution to EGI user communities engagement and support.
- Documentation made available to the community.
- Interoperability across multiple cloud management platforms.

Cycle #3, Sep 2012 – Mar 2013: Integration

- Focus on dev tools for management interfaces and clients for the test bed.
- Integration of the test bed services into the EGI infrastructure.
- Cloud brokering evaluation and deployment.
- Focus on use cases coordination and implementation.
- Opening of the test bed to early adopters.

Thank you.

Matteo Turilli

Senior Research Associate, OeRC, University of Oxford
Chair – EGI Federated Clouds Task Force
matteo.turilli@oerc.ox.ac.uk

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