



One year of The EGI Federated Clouds Task Force

HEPiX Fall 2012 Workshop – Beijing, 15-19 Oct

Matteo Turilli

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

Ian Collier

Grid Services Team Leader
Scientific Computing Department
STFC Rutherford Appleton Laboratory

Outline

- TF objectives, deliverables, mandate and membership.
- 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.

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, now extended to the end of EGI-InSPIRE 2014.

Activities: 3 blocks of 6 months each.

1. **Setup:** Sep 2011 – Mar 2012.

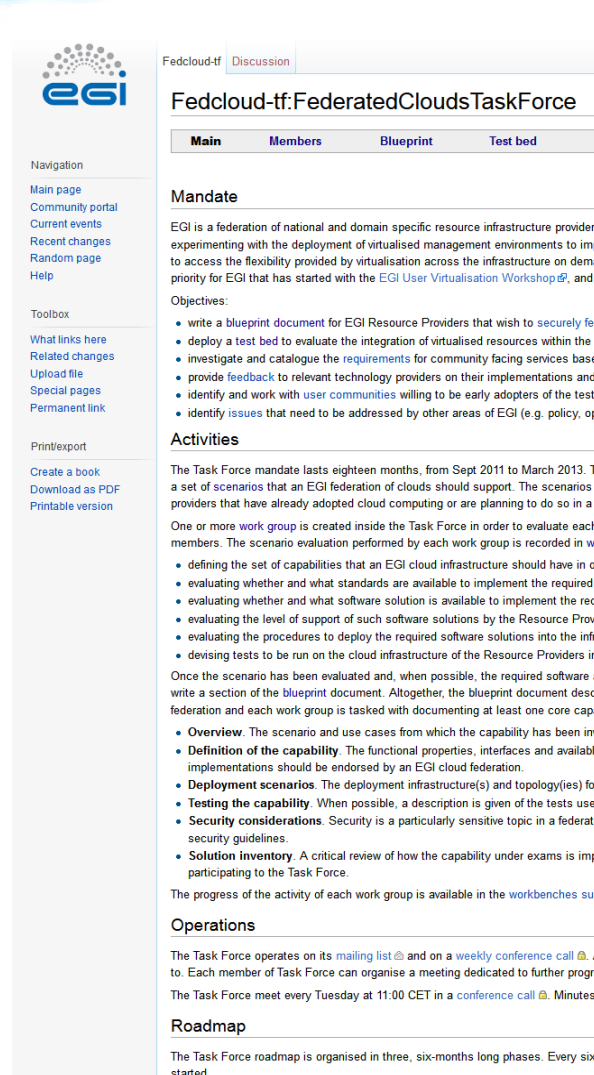
- Engagement of resource and technology providers.
- Federation model.
- Draft of the Blueprint document and demo.

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.



The screenshot shows the website for the Fedcloud-tf Federated Clouds Task Force. The page has a navigation menu on the left with links for Main page, Community portal, Current events, Recent changes, Random page, and Help. Below that is a Toolbox section with links for What links here, Related changes, Upload file, Special pages, and Permanent link. Further down are Print/export options and links to Create a book, Download as PDF, and Printable version.

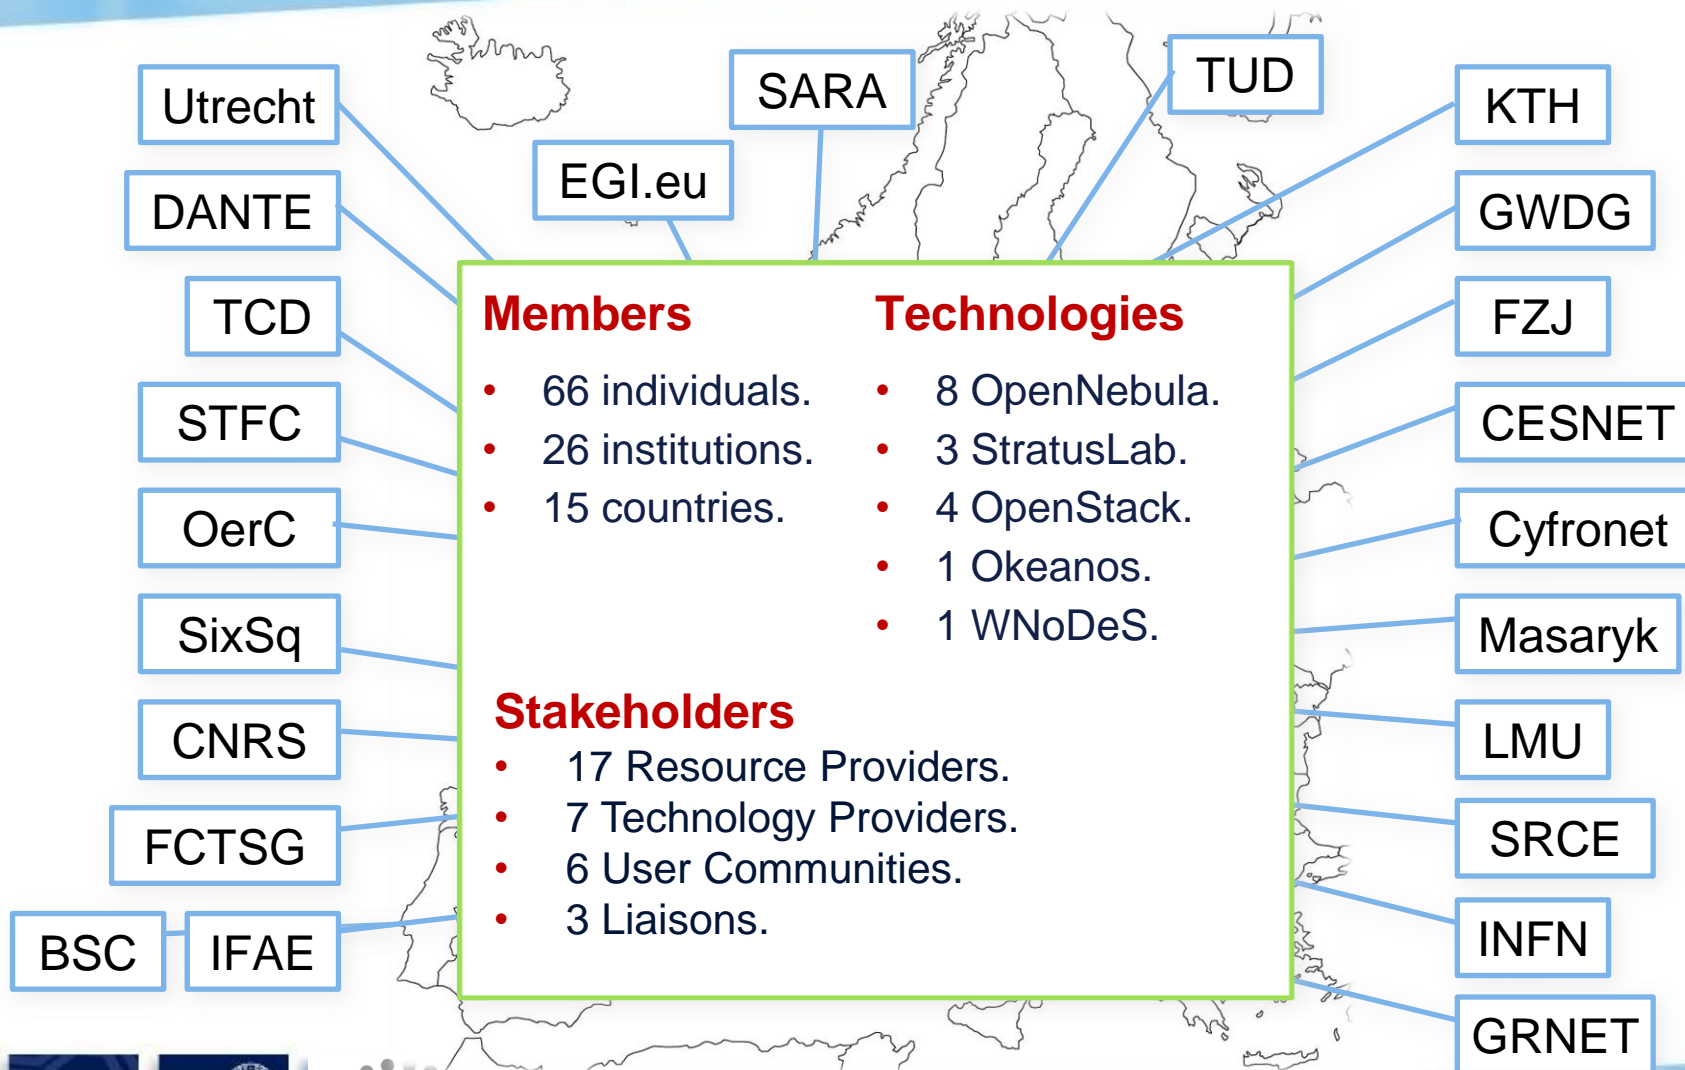
The main content area is titled "Fedcloud-tf: Federated Clouds Task Force" and has a sub-navigation menu with "Main", "Members", "Blueprint", and "Test bed". The "Mandate" section is highlighted, containing the following text:

EGI is a federation of national and domain specific resource infrastructure providers experimenting with the deployment of virtualised management environments to improve access to the flexibility provided by virtualisation across the infrastructure on demand priority for EGI that has started with the [EGI User Virtualisation Workshop](#), and the following objectives:

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

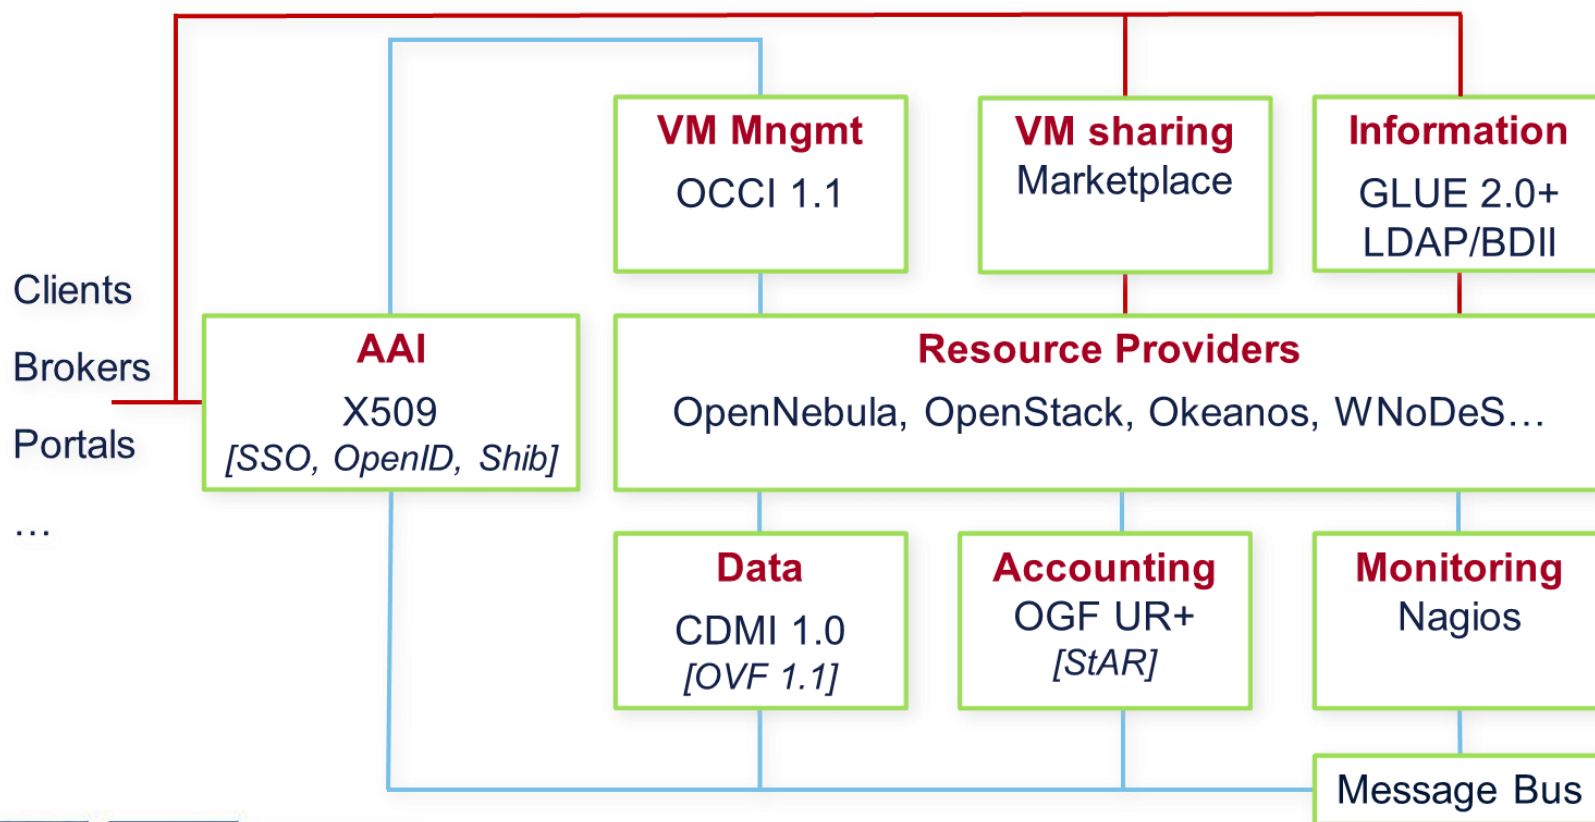
The "Activities" section follows, detailing the Task Force mandate from September 2011 to March 2013. It describes the creation of scenarios for an EGI federation of clouds, the formation of work groups, and the evaluation of capabilities and standards. It also lists key activities such as defining capabilities, evaluating standards, and documenting core capabilities. The "Operations" section mentions the Task Force's mailing list and weekly conference calls. The "Roadmap" section states that the Task Force roadmap is organized in three, six-month long phases.

Task Force Members and Technologies



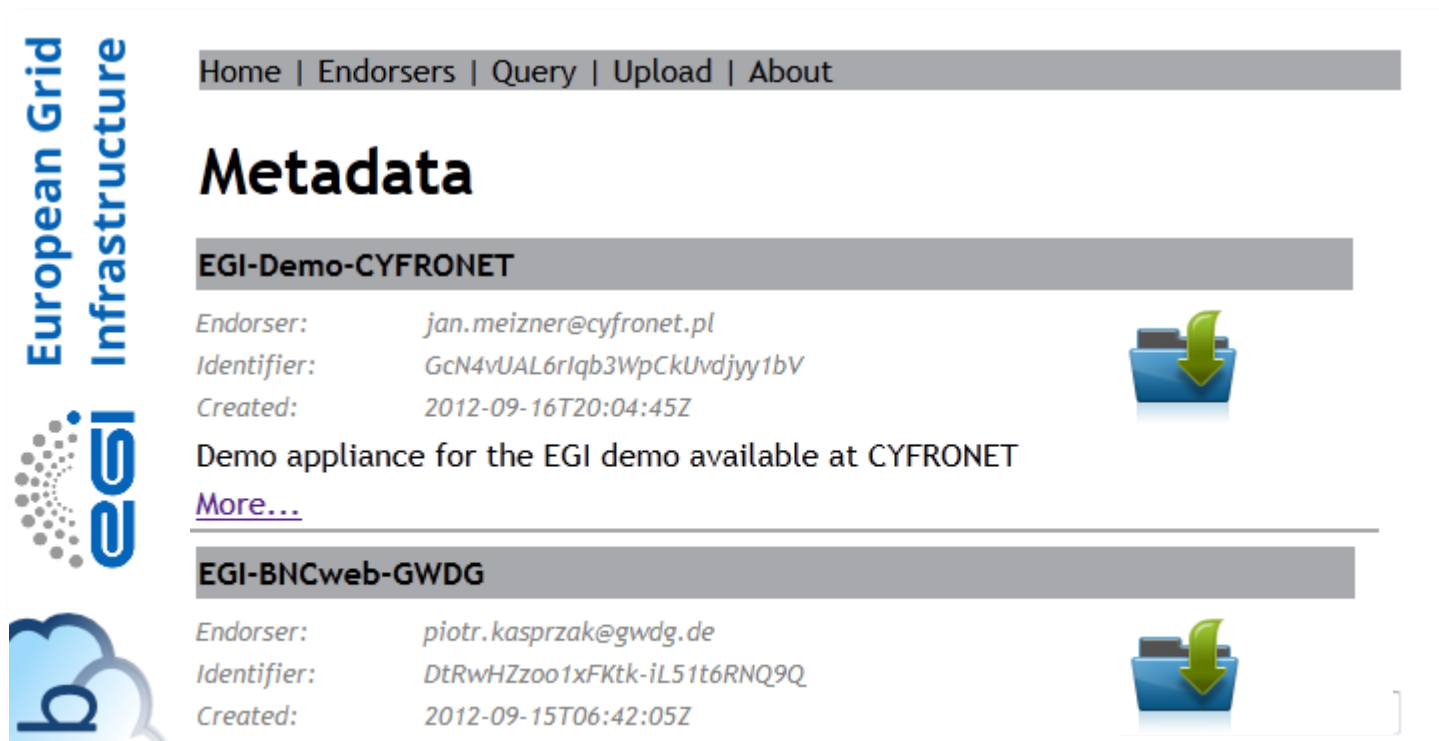
Federation Test bed – Oct 2012

Composed of 4 services, 2 management interfaces, 9 cloud infrastructures operated by 7 Resource Providers. 4 more providers are in the process of being federated.



Services

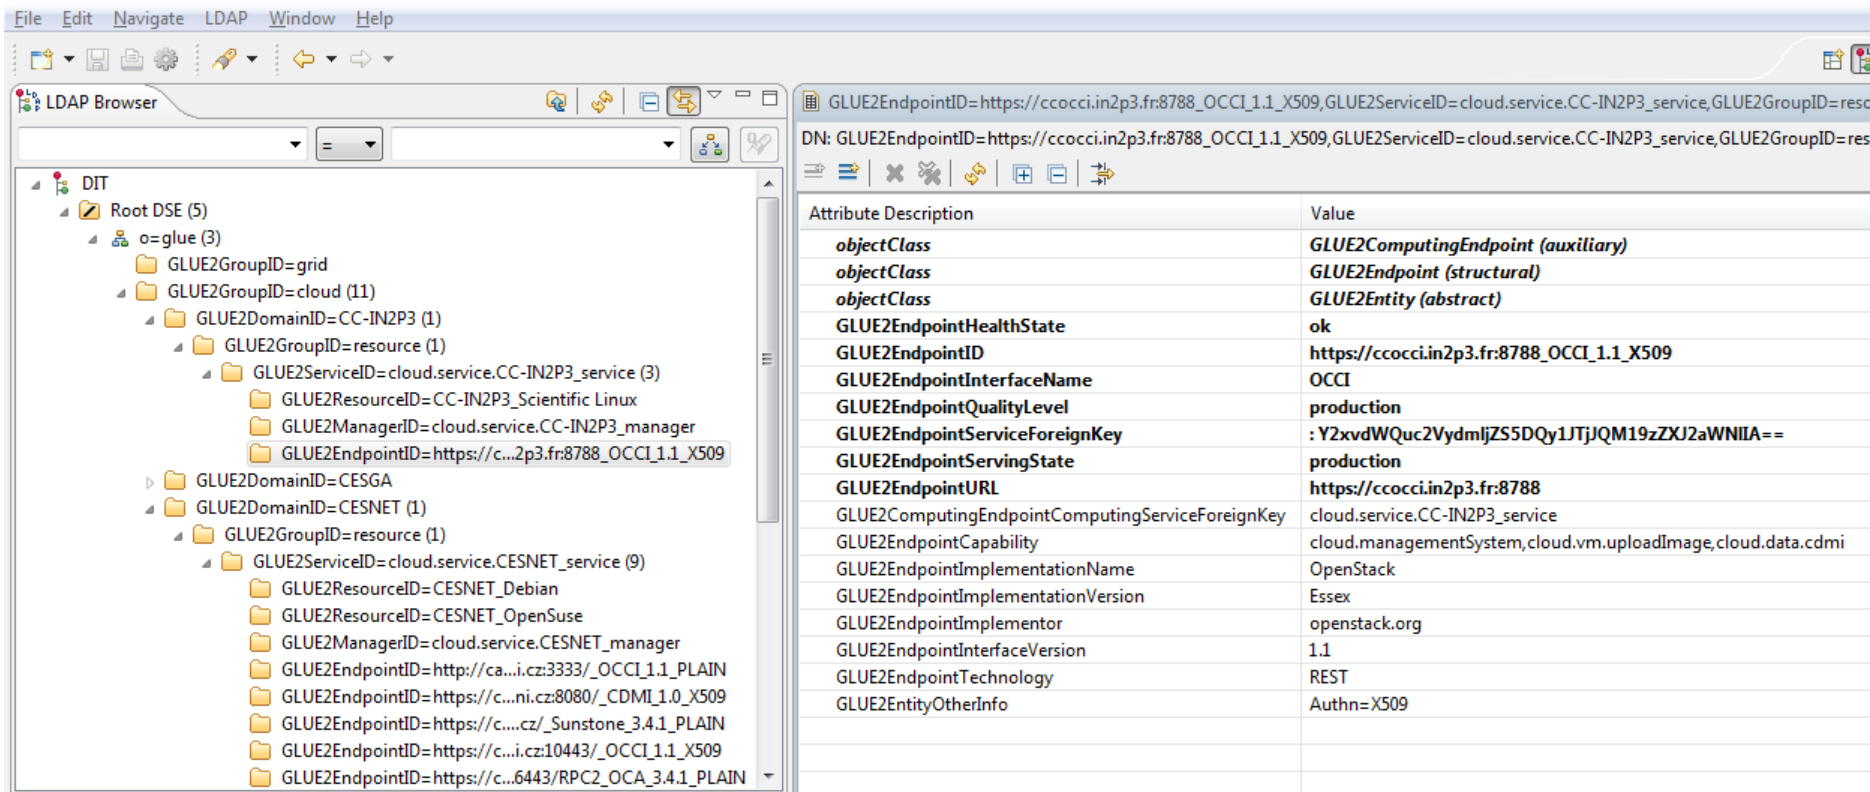
Marketplace. A repository where Resource Providers and EGI can publish metadata about images from which virtual machines can be instantiated. When needed, a single image can be signed and then endorsed by multiple providers.



The screenshot shows the EGI Marketplace Metadata page. On the left, there is a vertical navigation bar with the text "European Grid Infrastructure" and the "egi" logo. Below the logo is a blue speech bubble icon with a white letter 'b'. The main content area has a navigation bar with links: "Home | Endorsers | Query | Upload | About". The title "Metadata" is prominently displayed. Below this, there are two entries for metadata. The first entry is titled "EGI-Demo-CYFRONET" and includes the following information: Endorser: jan.meizner@cyfronet.pl, Identifier: GcN4vUAL6rlqb3WpCkUvdjyy1bV, and Created: 2012-09-16T20:04:45Z. To the right of this entry is a blue folder icon with a green arrow pointing down. Below the identifier is the text "Demo appliance for the EGI demo available at CYFRONET" and a link "More...". The second entry is titled "EGI-BNCweb-GWDG" and includes: Endorser: piotr.kasprzak@gwdg.de, Identifier: DtRwHZzoo1xFKtk-iL51t6RNQ9Q, and Created: 2012-09-15T06:42:05Z. To the right of this entry is another blue folder icon with a green arrow pointing down.

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.



The screenshot shows an LDAP Browser window with a tree view on the left and a detailed view on the right. The tree view shows a hierarchy starting with 'DIT' and 'Root DSE (5)', leading to 'o=glue (3)', 'GLUE2GroupID=grid', 'GLUE2GroupID=cloud (11)', 'GLUE2DomainID=CC-IN2P3 (1)', 'GLUE2GroupID=resource (1)', 'GLUE2ServiceID=cloud.service.CC-IN2P3_service (3)', 'GLUE2ResourceID=CC-IN2P3_Scientific Linux', 'GLUE2ManagerID=cloud.service.CC-IN2P3_manager', and 'GLUE2EndpointID=https://c...2p3.fr:8788_OCCE1.1.1_X509'. The detailed view on the right shows the following attributes and values:
















































Attribute	Description	Value
<i>objectClass</i>		<i>GLUE2ComputingEndpoint (auxiliary)</i>
<i>objectClass</i>		<i>GLUE2Endpoint (structural)</i>
<i>objectClass</i>		<i>GLUE2Entity (abstract)</i>
GLUE2EndpointHealthState		ok
GLUE2EndpointID		https://ccocci.in2p3.fr:8788_OCCE1.1.1_X509
GLUE2EndpointInterfaceName		OCCE
GLUE2EndpointQualityLevel		production
GLUE2EndpointServiceForeignKey		: Y2xvdWQuc2VydmljZS5DQy1JTjJQM19zZXJ2aWNIIA==
GLUE2EndpointServingState		production
GLUE2EndpointURL		https://ccocci.in2p3.fr:8788
GLUE2ComputingEndpointComputingServiceForeignKey		cloud.service.CC-IN2P3_service
GLUE2EndpointCapability		cloud.managementSystem,cloud.vm.uploadImage,cloud.data.cdmi
GLUE2EndpointImplementationName		OpenStack
GLUE2EndpointImplementationVersion		Essex
GLUE2EndpointImplementor		openstack.org
GLUE2EndpointInterfaceVersion		1.1
GLUE2EndpointTechnology		REST
GLUE2EntityOtherInfo		Authn=X509

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)				BDII Tests (BDII)				CDMI Tests (CDMI)			
Host	Status	Services	Actions	Host	Status	Services	Actions	Host	Status	Services	Actions
carach5.ics.muni.cz	UP	1 OK	 	cagnode42.cs.tcd.ie	UP	1 OK	 	bscgrid05.bsc.es	UP	1 OK	  
ccnovaapi.in2p3.fr	UP	1 OK	 	carach5.ics.muni.cz	UP	1 OK	 	carach3.ics.muni.cz	UP	1 OK	  
cloud-lab.grid.cyf-kr.edu.pl	UP	1 OK	 	cccldbdii01.in2p3.fr	UP	1 OK	 	cdmi.pdc2.pdc.kth.se	UP	1 OK	  
egi-cloud.zam.kfa-juelich.de	UP	1 OK	 	cloud-lab.grid.cyf-kr.edu.pl	UP	1 OK	 	occi.cloud.gwdg.de	UP	2 OK	  
front.redcloud.pdc.kth.se	UP	1 OK	 	egi-cloud.zam.kfa-juelich.de	UP	1 OK	 	Marketplace Tests (Marketplace)			
meghacloud.cesga.es	UP	1 OK	 	front.redcloud.pdc.kth.se	UP	1 OK	 	Host	Status	Services	Actions
occi.cloud.gwdg.de	UP	1 OK	 	occi.cloud.gwdg.de	UP	1 OK	 	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	 

Services

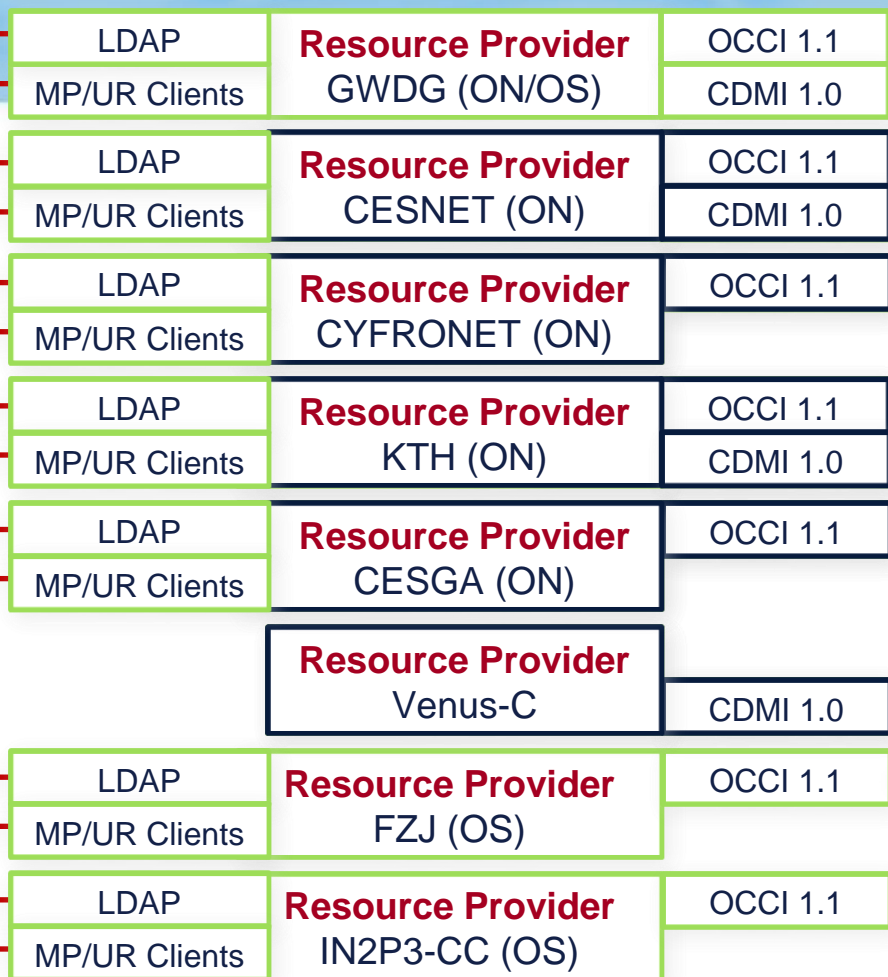
Accounting. Each cloud infrastructure generates usage records based on an extended version of the EGI UR format recommendation. Records are uploaded to a central server by means of a client customised for each type of infrastructure.

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

Federation Demo – Sep 2012



Information
GLUE 2.0
BDII

Monitoring
Nagios

VM metadata
Marketplace

Accounting
OGF UR
UR+ & StAR

ON = OpenNebula.
OS = OpenStack.
MP = Marketplace.
UR = Usage Records.

Message Bus

Blueprint Document



Fedcloud-tf [Discussion](#)

[Read](#)

[Edit](#)

[View history](#)

[Go](#)

[Search](#)

Fedcloud-tf:Blueprint:Introduction

(Redirected from [Fedcloud-tf:Blueprint](#))

[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 collection 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 interoperability across individual Cloud providers.

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

Contents [\[hide\]](#)

- 1 Introduction
- 2 High-level scenarios
 - 2.1 Scenario 1: VM Management
 - 2.2 Scenario 2: Managing my own data
 - 2.3 Scenario 3: Integrating multiple resource providers
 - 2.4 Scenario 4: Accounting across Resource Providers
 - 2.5 Scenario 5: Reliability/Availability of Resource Providers
 - 2.6 Scenario 6: VM/Resource state change notification
 - 2.7 Scenario 7: AA across Resource Providers
 - 2.8 Scenario 8: VM images across Resource Providers
- 3 Key Capabilities
 - 3.1 VM management
 - 3.2 Data management
 - 3.3 Information discovery
 - 3.4 Accounting

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



Blueprint Document

How to 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.
- **Space science** – ASTRA-GAIA project: data integration with scalable workflows.
- **Software Engineering**
 - SCI-BUS project: simulated environments for portal testing.
 - DIRAC project: framework for building ready to use distributed computing systems.



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

Ian Collier

Grid Services Team Leader
Scientific Computing Department
STFC Rutherford Appleton Laboratory

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>