EGI-InSPIRE

Integrating Resources into the EGI Production Infrastructure

EU MILESTONE: MS414

Document identifier:	EGI-MS414-V0-2.doc
Date:	**23/09/2010**
Activity:	**SA1**
Lead Partner:	**KTH**
Document Status:	**DRAFT**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/650

Abstract

<< The abstract should provide a brief neutral overview of the document and its contents and main conclusions. Once complete the abstract should be copied into the abstract field on the document server.>>This document describes and defines the operational interfaces that must be supported for resources to be integrated into EGI. This includes operational tools provided by the EGI-InSPIRE JRA1 activity and procedures and policies defined to ensure interoperability within EGI and in the interaction with other DCIs, the adoption of best practices and compliance with service level agreements.

<<document handling and production procedure is provided in https://documents.egi.eu/document/33 >>

I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE ("European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe") is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit http://creativecommons.org/licenses/by-nc/3.0/ or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: "Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration". Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Michaela Barth	KTH/SA1	
Reviewed by	**Moderator:** **Reviewers:** <<To be completed by project office on submission to AMB/PMB>>		
Approved by	**AMB & PMB** <<To be completed by project office on submission to EC>>		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
0	05/07/2011	Incomplete placeholder	Michaela Barth /KTH
1	06/07/2011	ToC	Michaela Barth /KTH
2	12/07/2011	Input on Operations Portal	Cyril L'orphelin / IN2P3

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE "Document Management Procedure" will be followed: https://wiki.egi.eu/wiki/Procedures

VI. TERMINOLOGY

A complete project glossary is provided at the following page: http://www.egi.eu/about/glossary/.

<<The authors should check if the acronyms are covered by the glossary page and if the definition is still correct; all the amendments should be communicated to glossary@egi.eu>>

VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting 'grids' of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today's production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

VIII. EXECUTIVE SUMMARY

<< The text should provide a summary of the full report so that the reader can 'in a page' understand the problem it has been written to cover. This includes an overview of the background material and motivation for the report, a summary of the analysis, and the report's main conclusions.>>

TABLE OF CONTENTS

1 INTRODUCTION

<< The ‘introduction’ of the document provides information on why it has been written, who the target audience is and what they will learn from reading it.>>

2 INTEGRATION OF MIDDLEWARE ON OPERATIONAL TOOL LEVEL

2.1 Interoperation at an Infrastructure Level

2.2 Overview Status of Middleware Integration for each Operational Tool

2.3 Definition and Description of a Management Interface

2.3.1 Functionality

2.3.2 Requirements

2.3.3 Integration of new Resources into GOCDB

Link to list of servicetypes on Wiki!

For bugs: https://savannah.cern.ch/bugs/?group=gocdb
New dev requests: RT
Operational/support issues: ggus.

See the list of external links at the bottom of the following page:
https://wiki.egi.eu/wiki/GOCDB/Documentation_Index

2.3.3.1 Integration of new MW service types

2.3.3.2 Integration of new non-MW service types

2.3.3.3 Declaration of new resources of an already available service type

2.3.3.4 Regular review of the list of available service types

2.3.3.5 Summary

2.3.3.6 Integration of gLite resources

No longer the whole list

2.3.3.7 Integration of ARC resources

As of release 0.8 of ARC, the ARC-CE runs a resource BDII with GLUE

schema 1.3, in the same way as gLite resources. Hence setting up a special site BDII is no

longer needed. More details are found in [R 22].
→ Verify

2.3.3.8 Integration of UNICORE resources

List as referred to in https://rt.egi.eu/rt/Ticket/Display.html?id=944

2.3.3.9 Integration of Globus resources

Update with latest discussion in Globus integration task force

2.4 Definition and Description of a Monitoring Interface

2.4.1 Functionality

2.4.2 Requirements

2.4.3 Interoperability of different MW Stacks with SAM/Nagios

2.4.4 Procedures to integrate new Nagios Probes

https://wiki.egi.eu/wiki/PROC07
https://wiki.egi.eu/wiki/PROC06

2.4.4.1 Tests and Nagios probes for gLite resources

2.4.4.2 Tests and Nagios probes for ARC resources

2.4.4.3 Tests and Nagios probes for UNICORE resources

2.4.4.4 Tests and Nagios probes for Globus resources

2.5 Definition and Description of an Accounting Interface

2.5.1 Functionality

2.5.2 Requirements

2.5.3 Current Status

Needs surely reference to EMI Compute Accounting working group

2.6 Definition and Description of a Support Interface

2.6.1 Functionality

2.6.2 Requirements

2.6.3 Integration of new Resources into GGUS

2.6.3.1 Integration of a new Resource Centre into the infrastructure

2.6.3.2 Integration of a new NGI into the infrastructure

2.6.3.3 Integration of a new Technology Provider into the infrastructure

2.7 Definition and Description of a Dashboard Interface

2.7.1 Functionality

In order to operate a distributed infrastructure, management and monitoring information has to be collected and presented in a labour saving way to assist the operators of the infrastructure in their daily work. The dashboard interface combines and harmonizes different static and dynamic information and therewith enables the operators to react on alarms, to interact with the sites, to provide 1st line support and/or to really operate the sites by creating and supervising problem tickets on regional as well as central level.

The dashboard allows predefined communication templates and is adaptable to different operational roles (1st line support, regional, central). Sites in the dashboard scope can be regional, central or predefined out of a list and can be sorted and displayed according to numerous criteria to indicate actions needed for a single service, but also for a whole region or even the whole production infrastructure.

2.7.2 Requirements

A dashboard interface has to fulfil the functionality described above.

With the increasing relevance of the SAGA Service Discovery specification [here] (OGF) for a standards-based approach for interoperability one more requirement on the dashboard is to provide such a well defined interface in order to be prepared for the harmonized integration of many different third party information providers.

We assume that EGI as a whole should try to unify the input:

- All sites should publish their information via a harmonized information service independently of the middleware stack used (e.g. GLUE2 based BDII)
- Access should be regulated by a harmonized user authentication service like VOMS or something better (see also detailed discussion in section 2.8).

Thus the dashboard and other tools don't have to be adapted to too many different information and authentication services.

In reality, though, it might be equally important to more directly connect to prevalent third-party information providers. A dashboard design that can effectively handle commonly used information services, especially those already established within EGI, while at the same time providing a well defined standard interface for interactions is the preferred solution.

2.7.3 The Operations Portal

The Operations Portal [here] content is based on information which is retrieved from several different distributed static and dynamic sources – databases, Grid Information System, web services, etc. – and gathered onto the portal. Interlacing this information has enabled us to display relevant views of static and dynamic information of the EGI production grid.

Integrating different technologies and different resources creates high dependencies to the data provided. Consequently, our technical solution is organized around a web service implementation that provides a transparent integration of each of these resources. The web service in question is named Lavoisier [here].

The goals of Lavosier are to provide:

- a web layer as independent as possible from the mechanisms technology used to retrieve the original information,
- intermediate information usable in the same format in order to cross-query it and
- information which is independent from the availability of the data provider.

This solution design means that the web application does not need to know the exact location of the data provider and neither which kind of technology has provided the information initially. All these concerns are already taken into account by Lavoisier.

Lavoisier has been developed in order to reduce the complexity induced by the various technologies, protocols and data formats used by its data sources. It is an extensible service for providing a unified view of data collected from multiple heterogeneous data sources. It enables us to easily and efficiently execute cross data sources queries, independently of used technologies. Data views are represented as XML documents and the query language is XSL.

The global architecture of the Operations Portal is presented in Fig. 1.

By using a plug-in schema, information can be retrieved from heterogeneous data providers (on the left side of the schema in Fig. 1). These plug-ins transform information in various formats extracted

from different technologies (i.e. RDMS, JSON, JMS, ldap, http, web service) into a standard format XML. At this stage it is easy to execute cross data sources queries by using XSLT transformation. In the end the web application is using all information in the same format (XML).

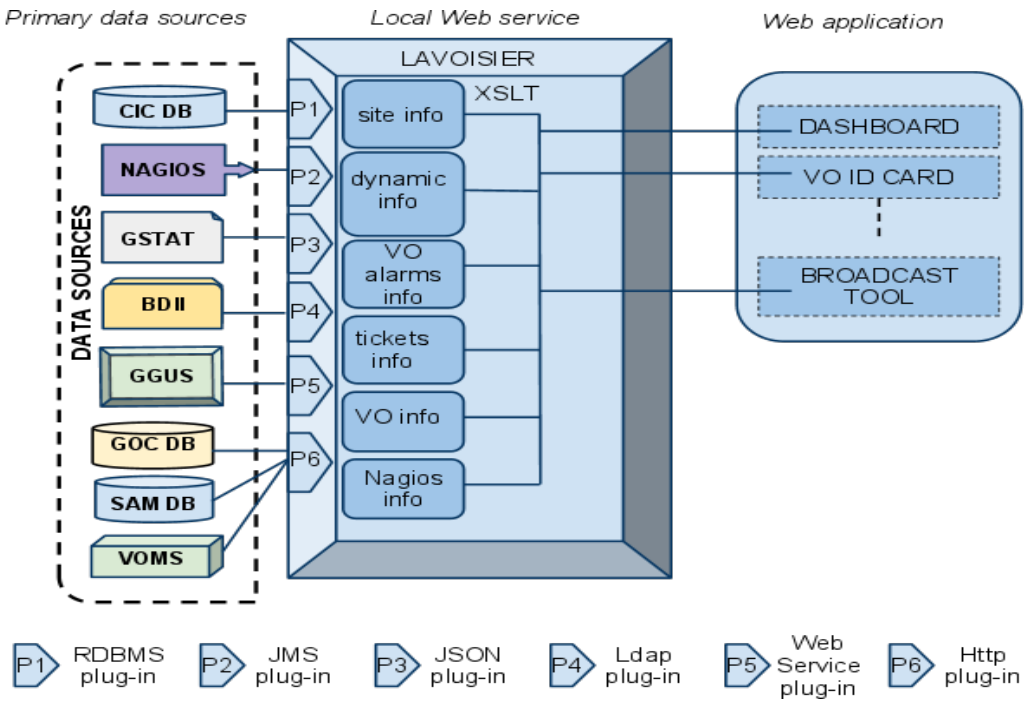

Fig. 1: Global architecture of the Operations Portal.

2.7.3.1 Integration of a new resource

The architecture of the portal has been designed to propose a standard access to information from an extended number of data sources. The integration of new data sources is eased by the use of the Lavoisier web service.

In the case of a known technology we will create and add a new view by using an existing plug-in out of the wide-range of plug-ins already available.

If a site and its resources are already integrated in all the other operational tools through existing information providers (e.g. registered in GOCDB, monitored by Nagios, publishing their information via BDII and having a tree in GGUS), existing plug-ins can be reused and no additional integration effort for the usage of the Operations Portal is needed.

For new providers, we will develop new plug-ins to be able to retrieve information from a new provider.

The integration of different information systems present in different middlewares such as ARC, UNICORE, or Globus can be done via an abstraction layer.

One such a possible abstraction layer could be to integrate the SAGA Service Discovery specification [here] (OGF) into a Lavoisier plug-in which will permit to access information using different services (like the information service of UNICORE – CIS [here]) and different schemas like CIM [here] or GLUE Schema [here] standards.

Lavoisier's flexibility allows us to be ready to integrate almost any kind of new information. Such an integration is certainly needed and meaningful for the new resource types entering EGI, such as HPC systems, virtualized resources or desktop resources. As long as these resources are monitored, it is possible to integrate them via plug-ins inside Lavoisier.

The integration will be done step-by-step during the whole project. The difficulty will be to identify the priorities in the components to integrate.

2.7.3.2 Alternative possibilities to integrate new information providers

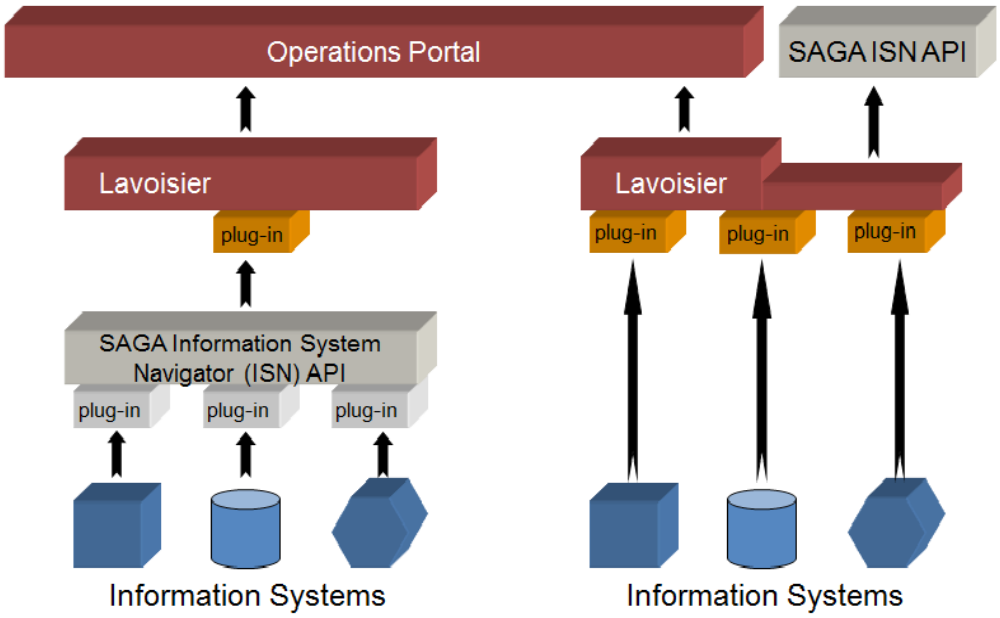

Fig. 2: Integration of new information systems into the Operations Portal

So far, no clear recommendation has been given yet on how to best include new information providers to the dashboard developers. The alternative depicted on the left side of the picture above might seem more work at first, but part of this work could probably be outsourced to the information providers and reused for other purposes. On the other hand, a Lavoisier to SAGA Information System Navigator (ISN) link might be needed anyway. The two possible alternatives are not mutually exclusive and might be combined.

2.7.3.3 Integration of a gLite resources

Plug-ins for all relevant information providers in the case of a site's gLite resources (Nagios, GOCDB, GGUS, BDII) exist and gLite resources can therefore be operated from within the Operations Portal.

2.7.3.4 Integration of a ARC resources

Plug-ins for all relevant information providers in the case of a site's ARC resources (Nagios, GOCDB, GGUS, BDII) exist and gLite resources can therefore be operated from within the Operations Portal.

2.7.3.5 Integration of a UNICORE resources

The UNICORE resources are registered in GOCDB and starting to be monitored by SAM/Nagios, GGUS trees exist. Hardware GLUE information could be taken from the Central Information Service CIS over the SAGA
ISN API link.

2.7.3.6 Integration of a Globus resources

Globus GT5 resources are registered in GOCDB and starting to be monitored by SAM/Nagios, GGUS trees exist.

Taking into account that LCG-CE is very similar to Globus GRAM, lcg-ce information providers can be reused for the BDII. With that Globus resources should be able to be directly integrated into the operational dashboard.

2.8 User Management, Authentication and Authorization

2.8.1 Functionality

2.8.2 Requirements

2.8.3 Argus

EMI has selected the ARGUS authorization framework as general approach for user

authorization based on the common SAML profile which shall be supported over all

middleware stacks.

2.8.3.1 Argus and gLite

2.8.3.2 Argus and ARC

2.8.3.3 Argus and UNICORE

2.8.3.4 Argus and Globus

3 INTEROPERATION AT PROCEDURES AND POLICY LEVEL

3.1 Scope

3.2 Current EGI Procedures and Policies

3.3 Future Procedures

4 OUTLOOK AND FUTURE PLANS

4.1 Operational requirements coming from NGIs

4.2 Operational requirements coming from Collaborations with other DCIs

5 REFERENCES

R 1	Operations Portal Home Page https://operations-portal.in2p3.fr
R 2	Lavoisier Home page http://grid.in2p3.fr/lavoisier
R 3	SAGA Service Discovery API http://www.ggf.org/documents/GFD.144.pdf
R 4	Common Information Service (CIS) for UNICORE Grids http://www.unicore.eu/community/development/CIS/cis.php http://www.d-grid.de/fileadmin/user_upload/documents/MonitoringWorkshop/Memon.pdf
R 5	Common Information Model Home Page http://www.dmtf.org/standards/cim/
R 6	GLUE schema http://infnforge.cnaf.infn.it/glueinfomodel/ Glue Schema specifications http://www.ogf.org/documents/GFD.147.pdf