EGI-InSPIRE

INTEGRATING RESOURCES INTO THE EGI PRODUCTION INFRASTRUCTURE

EU DELIVERABLE: MS407

Document identifier:	EGI-MS407-V0-4.odt
Date:	**19/08/2010**
Activity:	**SA1**
Lead Partner:	**KTH**
Document Status:	**DRAFT**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/111

Abstract

This document describes and defines the operational interfaces that must be supported for resources to be integrated into the EGI production infrastructure. This includes operational tools provided by activity EGI-JRA1 and procedures and policies defined together by O6, OE-13 and OE-11.

Delivery Slip

	Name	Partner/Activity	Date
From	Michaela Lechner	KTH/SA1	
Reviewed by	**Moderator:** **Reviewers:**		
Approved by	**AMB & PMB**		

Document Log

Issue	Date	Comment	Author/Partner
0.0	13/07/2010	Incomplete Placeholder	Michaela Lechner / KTH
0.1	30/07/2010	ToC after input from kickoff meeting	Michaela Lechner / KTH
0.2	05/08/2010	Added input on GOCDB	Gilles Mathieu / RAL-STFC
0.3	06/08/2010	Input on ARC and UNICORE	Michaela Lechner /KTH, Rebecca Breu / FZJ and Michael Grønager / DSMU
0.4	19/08/2010	Input Operations Portal, Globus and more UNICORE	Cyril L'orphelin /IN2P3 Rebecca Breu /FZJ Anton Frank /LRZ
1.0		First draft	
2.0			
4.0			
5.0			

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 the 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 the ESFRI projects. 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 − structured international user communities − that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

TABLE OF CONTENTS

1. INTRODUCTION....7

1.1. Purpose....7
1.2. Application area....7
1.3. References....7
1.4. Document amendment procedure....9
1.5. Terminology....9

2. EXECUTIVE SUMMARY....10

3. INTEGRATION AT OPERATIONAL TOOL LEVEL....10

3.1. Overview Interoperation Status for Operational Tools - MW....10
3.2. Interoperation at Infrastructure Level....11
3.2.1. Integration of new resources into GOCDB....11
3.2.2. Interoperability of different middleware stacks with Nagios....15
3.2.3. Operational Dashboard Portal....16
3.2.4. Accounting....18
3.2.5. GGUS....19
3.3. User management, Authentication and Authorisation....19
3.3.1. User management in gLite and ARC....19
3.3.2. User management in UNICORE....20
3.3.3. User management in Globus....20
3.4. Interoperation between operational Stacks....20
3.4.1. Job Submission....20
3.4.2. Data Management and Storage Infrastructure....20
3.4.3. Logging....20
3.5. Requirement lists to the middleware Providers....21
3.5.1. gLite....21
3.5.2. ARC....21
3.5.3. UNICORE....21
3.5.4. Globus....21

4. INTEROPERATION AT PROCEDURES AND POLICY LEVEL....21

4.1. Communication channels....21
4.2. Operation Level Agreements (OLA)....22
4.3. Procedure for adding a new NGI....22
4.4. Security....22
4.4.1. The integration into the EGI-CSIRT group....22

Index of Tables

Table 1: Table of references.....7
Table 2: Glossary of terms.....9
Table 3: Outlining the current status of interoperation for each MW stack relative to the operational tools.....11

Index of Figures

Fig. 1: Global architecture of the Operations Portal.....17

1. INTRODUCTION

1.1. PURPOSE

In order to add new resources into the EGI production infrastructure a basic set of operational interfaces that must be supported by the newcomers has to be defined.

Different resources will use different middleware components. EGI-InSPIRE will use the Unified Middleware Distribution (UMD). The UMD integrates middleware components provided by the European Middleware Initiative project (EMI), by the Initiative for Globus in Europe (IGE) project, and other external sources called “Community Contributions”. Services from the gLite, ARC and UNICORE middleware stacks will be included in the EMI release. Within the scope of this document middleware stacks collected in the UMD are taken into account.

Operational tools such as the GOC Database (GOCDB) or the Nagios monitoring tools, are key software components for a reliable and stable operation and monitoring of the infrastructure. The current set of as basic considered operational tools is inherited from our experiences within EGEEIII. However this might change in the future. Still we take this a starting point when comparing the interoperability of different middleware components for each operational tool in our current horizon.

Equally operational procedures and policies are needed to enforce the application of the agreed basic set of operational interfaces to be supported by all resources. Some old procedures and policies valid in EGEEIII can possibly be adapted, new, not-yet-existent but needed ones have to be identified and split down to requirements. Special focus shall be laid on security.

1.2. 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.

1.3. REFERENCES

Table 1: Table of references

R 1	MS405: Operational Security procedures https://documents.egi.eu/secure/ShowDocument?docid=47
R 2	EGI Wiki https://wiki.egi.eu/
R 3	EGI SSO: https://www.egi.eu/sso/
R 4	EGI Mail manager https://mailman.egi.eu/mailman/listinfo
R 5	GOCDB requests and wish list https://savannah.cern.ch/support/?group=gocdb
R 6	GOCDB general documentation index: http://goc.grid.sinica.edu.tw/gocwiki/GOCDB_Documentation_Index
R 7	dCache http://www.dcache.org/

R 8	LFC catalogue service http://goc.grid.sinica.edu.tw/gocwiki/How_to_set_up_an_LFC_service
R 9	VOMS http://hep-project-grid-scg.web.cern.ch/hep-project-grid-scg/voms.html
R 10	WLCG http://lcg.web.cern.ch/lcg/public/
R 11	Globus Meta Data Service, Globus MDS
R 12	M.Ellert et al., Future Generation Computer Systems 23 (2007) 219-240.
R 13	Field L and Schultz M W Proc. of CHEP 2004, CERN-2005-002, 2005
R 14	GLUE schema http://infnforge.cnaf.infn.it/glueinfomodel/ Glue Schema specifications http://www.ogf.org/documents/GFD.147.pdf
R 15	gLite WMS http://glite.web.cern.ch/glite/packages/R3.0/deployment/glite-WMS/glite-WMS.asp
R 16	EGEE Accounting Portal http://www3.egee.cesga.es/
R 17	Real Time Monitor http://gridportal.hep.ph.ic.ac.uk/rtm/
R 18	UNICORE bug tracker http://sourceforge.net/tracker/?group_id=102081&atid=633902 UNICORE feature tracker http://sourceforge.net/tracker/?group_id=102081&atid=633905
R 19	SGAS http://www.sgas.se
R 20	SGAS to APEL Byrom R et al. http://www.gridpp.ac.uk/abstracts/allhands2005/apel.pdf
R 21	Grønager M et al escience, pp.493-500, 2008 Fourth IEEE International Conference on eScience, 2008
R 22	Towards Sustainability: An Interoperability Outline for a Regional ARC based infrastructure in the WLCG and EGEE infrastructures
R 23	Operations Portal New Home Page https://operations-portal.in2p3.fr
R 24	Lavoisier Home page http://grid.in2p3.fr/lavoisier
R 25	SAGA Service Discovery API http://www.ggf.org/documents/GFD.144.pdf
R 26	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 27	Common Information Model Home Page http://www.dmtf.org/standards/cim/
R 28	UNICORE support mailing lists for EMI related and general issues: emi-support@unicore.eu and unicore-support@lists.sourceforge.net.

R 29	Google maps CIS web client demo http://omiiei.zam.kfa-juelich.de:6001/web/Index
R 30	UNICORE 6 Monitoring with Nagios http://www.d-grid.de/fileadmin/user_upload/documents/MonitoringWorkshop/Rambadt.pdf
R 31	PL-Grid UNICORE Monitoring System http://www.unicore.eu/summit/2010/presentations/18_Bala_Monitoring.pdf

1.4. 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

1.5. TERMINOLOGY

A complete project glossary is provided in the EGI-InSPIRE glossary:

http://www.egi.eu/results/glossary/.

The table below contains further terminology not provided in the previous location:

ARC	Advanced Resource Connector (middleware stack)
BDII	Berkeley Database Information Index
CIS	Common Information Service
GIIS	Grid Index Information Server
gLite	Lightweight Middleware for Grid Computing (middleware stack)
Globus	Globus Toolkit Grid Middleware (middleware stack)
GRIS	Grid Resource Information Service
GOCDB	Grid Operations Centre DataBase
UNICORE	Uniform Interface for Computing Resources (middleware stack)
XUUDB	UNICORE User Database

Table 2: Glossary of terms.

2. EXECUTIVE SUMMARY

This document describes and defines the operational interfaces that must be supported for resources to be integrated into the EGI production infrastructure.

For the each of the operational tools considered we provide a general impression of what has to be done to integrate a new middleware stack, followed by a detailed analysis for each middleware stack in the UMD on the status quo and the immediate future plans.

An overview table shows the general picture outlining the current status of each MW in relation to the operational tools.

Outgoing from that, we conclude a requirement list of our suggestions to each MW provider, so that sites running only this specific MW stack will still be able to make full use of all relevant operational tool features and therefore be fully integrated. Requirements can also steem from a more general interoperability point of view.

For completeness we also include a section touching the level of interoperation level between the operational stacks, since this will get more important in the future when sites start to support several middleware stacks in parallel.

Additionally this document will try to give an overview of the status of operational procedures and policies needed for the integration of new resources.

3. INTEGRATION AT OPERATIONAL TOOL LEVEL

EGI-InSPIRE scans eagerly the blueprint on how to successfully run a common European Grid as handed over by the end of EGEE III. A certain amount of optimization as well as decisiveness to pick out the pearls will be needed in the formidable quest to create a sustainable model for operating a growing pan European Grid infrastructure that builds on national and regional funded Grid initiatives who want to work together.

Availability and reliability measurement, registration of services, indexing, monitoring, accounting, user and operational support in the EGI Grid infrastructure currently rely on operational tools already developed in the framework of the EGEE project series. Tool development is an ongoing effort and is part of the EGI-InSPIRE JRA1 work programme [R2].

While different harmonized middleware stacks are supported by EGI for deployment in the resource centres, the central and distributed instances of the operational tools are operated by a small number of partners committed to provide such services for National or Regional Grid Initiatives, or even for the whole EGI.

The EGI infrastructure will have several middleware stacks deployed within it. Presently, as a result of the EGEE and WLCG projects, only gLite is fully integrated into all the operational tools, whilst ARC has been partially integrated, and for Globus and UNICORE operational integration is still to be implemented. Comprehensive integration is a short-term objective of the project.

In a second phase, it is expected that site administrators and user communities will provide requirements for the interoperability between different middleware stacks, and that the EGI infrastructure will be integrated with new types of resource, such as digital libraries and repositories, desktop grids, High Performance Computing, etc.

3.1. OVERVIEW INTEROPERATION STATUS FOR OPERATIONAL TOOLS - MW

	gLite	ARC	UNICORE	Globus
GOCDB	completed	completed	to be done	to be done
Nagios - Definition of critical tests	completed	completed	?	to be selected (available in NGI-DE, IGE)
Nagios - Probes	completed	completed	?	to be selected (available in NGI-DE, IGE)
Operational Dashboard	completed	completed	to be done	to be done
Accounting			not (yet) available	not (yet) available
3rd level support in GGUS				to be done

Table 3: Outlining the current status of interoperation for each MW stack relative to the operational tools

3.2. INTEROPERATION AT INFRASTRUCTURE LEVEL

An important operational interface of a resource is the capability to be put in downtime if under maintenance, the capability to undergo a certification process and thereby reach production status and the capability to be monitored to assess its operational security level. Within the current EGI production infrastructure GOCDB is the tool of choice for fulfilling these tasks. It portrays what services are running where and who to contact on a management and technical level as well as in case of security issues.

A first step towards integration of resources is therefore to enable the registration of the possibly new kind of services provided by these new resources in GOCDB.

The next step is to monitor the resources in some way using the OGSA GLUE2 standard schema enabling a unified view of Grids and resources per infrastructure, computing center or federation. One such a possible monitoring tool is for example Nagios, where critical services and probes for them have to be defined. The collected information can then be plugged in in the Operational Portal to give a detailed overview of operational status and the possibility to contact the sites as stored in GOCDB. General monitoring is also needed to produce Availability & Reliability figures.

Besides that we have to insure the quality of service by providing 3rd level support and by introducing accounting.

EMI promised to set up a 3rd level structure within GGUS for its various middleware stacks and services since GGUS proved to be a decent tool for that during the EGEE area.

Accounting is important as well since the key feature of an operational infrastructure is that the resources have high availability and reliability, and that we can measure their usage.

3.2.1. Integration of new resources into GOCDB

Resources are stored in GOCDB using the following two basic concepts:

1. "**Service types**", which represent generic components deployable on the Grid infrastructure. They can be middleware components (e.g. CE, WMS, SRM...) or components specific to the operational infrastructure (e.g. MessageBroker, RegionalNagios...).
2. "**Service endpoints**", which represent deployed instances of a service type.

In order for new resources to be integrated into GOCDB, the type of these resources has to be integrated first as "service type", and then the deployed instances of this service type can be declared.

Because of that model, integrating new resources to GOCDB does not require any development effort.

It is a matter of adding the proper set of information to the existing system as described in the following sections.

3.2.1.1. Integration of new MW service types

New MW service types can either be new services from an already listed middleware, or services from a new middleware stack.

In the first case, the proposed procedure is as follows:

EGI-JRA1 gets from middleware providers (EMI) the information about new services that have been added to an existing middleware stack.

In the second case, the request of adding a set of services belonging to a previously undeclared MW stack implies that a verification has been made about the validity of the request. This is to ensure that only MW stacks officially supported are actually integrated to GOCDB.

Request for adding the new service types in GOCDB should be made through the official request submission channel [R5].

A validation board described in 3.2.1.4 discusses the request and gives its green light for the integration. New service types are then added to GOCDB and are made available to declare new resources as described in 3.2.1.3.

3.2.1.2. Integration of new non-MW resources types

There is a need to store non-middleware service types in GOCDB since services used for Grid operations are declared within the repository. Also, there might be a need to store and present information about application services, deployed by Grid sites to support certain VOs without belonging to a specific middleware distribution.

New non-middleware services are integrated into GOCDB in a similar way to MW services, apart from the fact that the initial information doesn't come from EMI but from either EGI-JRA1 itself (in case of this being a service for operations management) or from a user community (in case of an application specific service).

The way to deal with the request and eventually integrate the new type to GOCDB is similar to what is described in section 3.2.1.1.

3.2.1.3. Declaration of new resources of an already available resource type in GOCDB

Once a service type is integrated into GOCDB, instances of this service can be declared as service endpoints. This is done by the resource providers (i.e. administrators of the site hosting the endpoint, regional managers, operational staff). A complete description of the process is described in the GOCDB user documentation [R 6].

3.2.1.4. Regular review of the list of available service types

The normal evolution of any infrastructure and middleware stack means that some service types might become obsolete with time. To avoid filling up GOCDB with too many disused services, a regular review of the list of available service types will be made. This task will be under the responsibility of GOCDB developers, who will get information from EMI and EGI-JRA1 before producing a list of service types candidate for decommission.

3.2.1.5. Summary of the complete procedure

The complete procedure to have new resources integrated to GOCDB is as follows:

- If the service type is already available in GOCDB, service endpoints can be added following [R 6] .
- If the service type is not available, a request to GOCDB developers has to be made in [R 5]. The case will then be discussed as described above and eventually result in the new service type being added.

3.2.1.6. Integrating gLite resources in GOCDB

Xxxxx (something new needed here?, or just statement that: Current resources are integrated. New resources are added according to the procedure as described above.

Possibly list of currently listed service types related to gLite???

CE, WMS, SRM...

)

3.2.1.7. Integrating ARC resources in GOCDB

ARC resources were first added into GOCDB already a long time ago. This has happened even though the Nordic infrastructure using the ARC middleware was not formally an EGEE partner. ARC integration could therefore serve as a role model on how to integrate other middleware stacks. In the beginning a lot of services were already common to gLite, like storage elements (dCache), catalogue service (LFC), VOMS, etc.

However the ARC method of dynamic service indexing, the ARC GIIS and the ARC-CE were not supported in GOCDB. The ARC-CE was added as a new Compute Element service type. A virtual site was created for NDGF and the ARC-CEs could be registered there.

For the indexing of services another solution was chosen. ARC had applied the Globus Meta Data Service consisting of top level GIIS and site level GRIS services. In order to be visible for EGEE services a special BDII has been set up for the virtual NDGF site which dynamically collected the content of the GRIS'es of the ARC-CEs based on the list of Ces provided by the GIIS'es. Recently the BDII concept has been adapted by the ARC-CE which of release 0.8 supports direct rendering to the GLUE schema. Hence setting up a special site BDII is no longer needed.

Nowadays new resources are simply added according to the procedure as described above.

3.2.1.8. Integrating UNICORE resources in GOCDB

The needed MW service types haven't been defined yet in GOCDB.

A list of service types that need to be defined follows. The different service types are typically installed on separate machines, but don't need to be.

- Gateway (a site's entry point)
- Registry
- Workflow Engine
- Service Orchestrator
- UNICORE/X
- CIS
- XUUDB
- Target System Interface (TSI) (submits jobs to the local batch system)

Some of these are quite tightly coupled, though, and are not visible as separate services to clients, nor can they be tested separately. Thus it might not make sense to separate them when integrating them.

3.2.1.9. Integrating Globus resources in GOCDB

The three most important service types for Globus which need to be registered into the GOCDB are:

1. job submission service for Globus version 4.0.x, 4.2.x (WS-GRAM) and 5.x (GRAM5).
2. data transfer (GridFTP) service.
3. certificate based interactive login service (gsisshd).

Used ports can differ from the default, thus the registration of the port must be possible as well.

3.2.2. Interoperability of different middleware stacks with Nagios

Xxxxx (Nagios as well-known, mature test execution and notification environment, defininig critical tests and writing probes)

Grid monitoring is needed to ensure the infrastructure's reliability and find causes of failures immediately. Statistical data is collected to provide input for the availability and reliability figures to see if OLAs are fulfilled and production level is reached. Users and operators are informed about the state of the Grid. Nagios is one possible input plugin for the Operational Portal.

3.2.2.1. Critical tests and Nagios probes for gLite resources

Xxxxx

3.2.2.2. Critical tests and Nagios probes for ARC resources

Historically Nagios' predator the Service Availability Monitoring, SAM, was the first EGEE infrastructure service to interact with ARC services. Every 3rd hour SAM executed tests against the different sites registered in the GOCDB by querying the individual services listed in the site BDII. SAM tests for index, storage, catalogue could run right from the start. A new sensor suite in the modular SAM was developed for the new Compute Element service type ARC-CE. The WLCG Management Board and an extra working group made sure that the tests for the different CE types compare and a fair and balanced translation between the different CE tests is ensured.

The transition towards Nagios monitoring was done during EGEE III together with gLite.

3.2.2.3. Critical tests and Nagios probes for UNICORE resources

UNICORE does have Nagios reporters which make use of the UNICORE monitoring tool SIMON. The Site Monitor for UNICORE resources (SIMON) [R 30] submits various kinds of UNICORE test jobs to check the availability of the UNICORE stack. One could integrate those into the EGI Nagios. SIMON acts as a user, thus needs its own certificate, login and entry in the UNICORE User Database. PL-Grid defined a number of critical tests and their dependencies [R 31].

UNICOREs Common Information Service (CIS) [R 26] provides detailed information about the underlying system, e.g. the number of CPUS, memory, number of running jobs etc. according to the OGSA standard GLUE2 information model [R 14] for representing resource information. A small demo of a Google maps CIS web client can be found under [R 29].

3.2.2.4. Critical tests and Nagios probes for Globus resources

Critical tests for Globus are the availability of the servers for central services (RFT, MyProxy, MDS/WEBMDS) and of the services at the resources (GSI-SSH, GridFTP, (WS-)GRAM, etc.).

Various Nagios probes have been developed in the scope of D-Grid/NGI-DE and DEISA.

Currently the following Nagios probes for critical tests are available:

- Globus service availability (GSI-SSH, GridFTP, (WS-)GRAM)

- GridFTP server availability test
- WS-GRAM (Globus v. 4.0.x) job submission test
- GridFTP file transfer test
- Globus container certificates (availability, lifetime)
- Globus container memory consumption
- RFT PostgreSQL DB
- RFT transfer test
- Globus WebMDS status
- Globus WebMDS HTTP response
- Version check of IGTF CA distribution
- Host certificate validity life-time check

It has to be checked if these Nagios probes can be used as is or if they need to be adjusted to the EGI requirements.

3.2.3. Operational Dashboard Portal

The Operations Portal [R 23] 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. Criss-crossing this information has enabled us to display relevant views of static and dynamic information of the EGEE, now 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 for a transparent integration of each of these resources. The web service in question is named Lavoisier [R 24].

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 implicates that the web application doesn't 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 on Fig. 1.

By using a plug-in schema we are able to retrieve information 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).

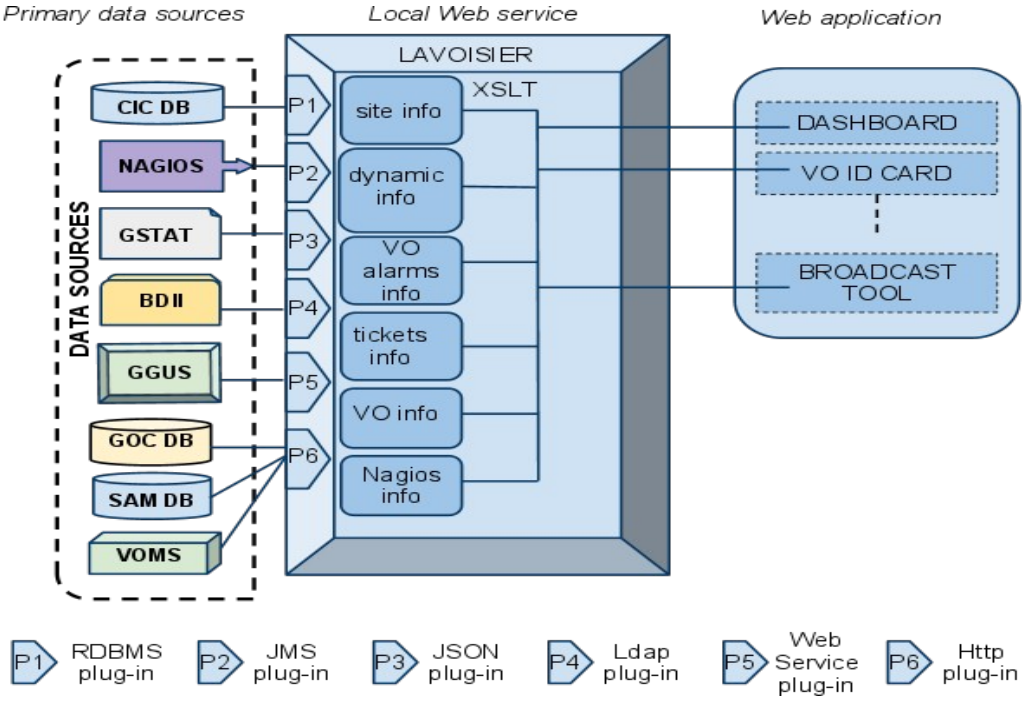

Fig. 1: Global architecture of the Operations Portal.

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

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 will be done via an abstraction layer.

One such a possible abstraction layer could be to integrate the SAGA Service Discovery specification [R 25] (OGF) into a Lavoisier plug-in which will permit to access information using different services (like the information service of UNICORE – CIS [R 26]) and different schemas like CIM [R 27] or Glue Schema [R 14] 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 coming into the EGI production infrastructure, such as HPC systems, virtualized resources or

desktop resources. As long as these resources are monitored we are able 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.

3.2.3.2. gLite resources in the Operational Dashboard

gLite resources are Nagios monitored and therefore already integrated. (true?)

3.2.3.3. ARC resources in the Operational Dashboard

ARC resources are Nagios monitored and therefore already integrated. (true?)

3.2.3.4. UNICORE resources in the Operational Dashboard

Xxxxx (depends on other things not finished yet, (list of services needed and integrated).

(Indirect over SIMON Nagios or directly from CIS over SAGA Link?)

3.2.3.5. Globus resources in the Operational Dashboard

With the Nagios probes available the operational alarms from Globus resources or central servers can be directly integrated in the operational dashboard.

3.2.4. Accounting

Xxxxx (general John Gordon)

(quotes from Tiziana:)

Accounting is a critical aspect as heavily relies on the adoption of a common usage record standard which doesn't exist at the moment.

Also, a common transport mechanism needs to be identified to transport records across sites deploying different middleware stacks.

Accounting of MPI jobs is also of great importance, together with accounting of virtual resources (grid-cloud integration).

3.2.4.1. gLite resources

Xxxxx

3.2.4.2. ARC resources

Accounting integration was performed already during EGEE III. The quest was to gather and export accounting from the Nordic T1 and T2s, which for the compute part were based on ARC, sorted per VO to the EGEE Accounting Portal. The EGEE Accounting Portal uses the APEL database as back-end, and direct DB insertion is provided per site. ARC-CE supports accounting via SGAS (SweGrid Accounting System) and an automatic script for exporting the accounting info gathered in SGAS to APEL was set up.

There is even an added value of this approach; Many national states do not allow for accounting info on the person level to be exported outside country borders. Hence a federated approach only submitting accounting info on a relevant level will be required for many states.

3.2.4.3. UNICORE resources

Currently no means of collecting accounting and usage records are directly implemented within UNICORE. Instead, this is done directly via the underlying batch system, see for example as in the DEISA project.

3.2.4.4. Globus resources

OGF-UR is available and used in DEISA (and soon also in PRACE). However, currently it is not integrated in the Globus tools. There were efforts of adopting DGAS for Globus in the scope of D-Grid. It was also planned to use OGF-UR there (which was unfortunately not yet provided by DGAS at that time).

3.2.5. GGUS

Xxxxx (3rd Level support, general Thorsten Antoni, ensuring that informations flows between different parts of the infrastructure)

3.2.5.1. gLite resources

Xxxxx

3.2.5.2. ARC resources

Xxxxx

3.2.5.3. UNICORE resources

UNICORE is part of EMI. EMI will soon set up a structure within GGUS for its various services, including UNICORE. Problems that can't be solved within EGI or EMI will be relayed to UNICORE's bug and feature tracker [R 18] or to the support mailing lists [R 28].

(needs: statement from EMI on that)

3.2.5.4. Globus resources

3rd level support for Globus will be provided by IGE. IGE is setting up a support infrastructure for the European Globus users in all European, national, and regional e-Infrastructures with EGI and DEISA/PRACE being the most important ones. GGUS will contain a queue to forward 3rd level support tickets directly to the IGE user support team. Further details will be clarified shortly after the project start of IGE.

3.3. USER MANAGEMENT, AUTHENTICATION AND AUTHORISATION

(new point suggested by Rebecca from UNICORE, should include references to VOMS, proxies, Central user databases and like in DEISA, LDAP server and generally how roles and certificates currently are matched to user accounts at the site.)

(ev. Refer to https://wiki.egi.eu/wiki/EGI_IGTF_Release_Process

)

3.3.1. User management in gLite and ARC

VOMS is used for VO and user administration.

3.3.2. User management in UNICORE

The UNICORE User Database (XUUDB) stores the mapping of user certificates/DN's to logins and roles at the sites. Proxy certificates are not used in UNICORE. Technically, it doesn't matter who manages the XUUDB user database. Every site can set up their own XUUDB and an independent way of managing it, or there could be a central XUUDB, or a central service that generates input for each site's XUUDB (which is, so far, the most common way).

3.3.3. User management in Globus

Globus first of all relies on the entries in the Globus grid-mapfile for authorization purposes. VOMS of VOMRS can be used to provide the necessary entries in order to achieve a high-level VO management for Globus.

3.4. INTEROPERATION BETWEEN OPERATIONAL STACKS

Requirements on the interoperation between different operational stacks have to come from the sites and NGIs which run more than one MW. Issues around the interoperability of the middleware itself from a user point of view remain out of scope for the context of this milestone, though.

3.4.1. Job Submission

Cross Grid job submission is not strictly needed for infrastructure interoperation. It does, however, promote a more seamless integration between different infrastructures and middleware stacks and might therefore be desirable.

3.4.1.1. Direct submission from gLite-WMS to an ARC-CE

Unlike gLite based Computing Elements (CEs), ARC-CEs are accessed directly by the ARC client. No intermediate resource broker like the gLite-WMS is needed. Brokering is performed by the ARC client. Various schemes have been explored for submitting a job from a gLite-based infrastructure to an ARC environment. Already in EGEE-II-SA3 [R] direct submission from gLite-WMS to ARC-CEs got implemented and is now part of the standard gLite-WMS.

3.4.1.2. Direct submission from ARC to a gLite-CREAM-CE

The latest (which one?) version of the ARC client also supports submission directly to the gLite-CREAM-CE hence making cross grid submission possible also in the ARC->gLite direction. A more detailed analysis on how ARC-CE compares to gLite flavoured CEs has been written in [R]

3.4.2. Data Management and Storage Infrastructure

(dCache for ARC and gLite)

3.4.3. Logging

Collecting detailed logging information in a common accessible and exchangeable format is clearly not the most essential high priority task when integrating new resources. A detailed analysis of this might follow in a later milestone.

3.4.3.1. Real Time Monitoring of ARC resources

ARC-CE supports logging calls to the gLite LB server. This means that detailed job states can also be obtained by ARC-CE sites, enabling advanced real time monitoring of the production flow like in the Real Time Monitor [R.] and for debugging scenarios.

Instead of installing the logging clients on all the ARC-CEs a hook in ARC for directly exporting the detailed job states to the gLite Logging and Bookkeeping server was chosen.

3.5. REQUIREMENT LISTS TO THE MIDDLEWARE PROVIDERS

Xxxxx (outgoing from points in 3.2)

(possibly related: https://rt.egi.eu/rt/Ticket/Display.html?id=231)

3.5.1. gLite

Xxxxx

3.5.2. ARC

Xxxxx

3.5.3. UNICORE

Xxxxx

3.5.4. Globus

Xxxxx

4. INTEROPERATION AT PROCEDURES AND POLICY LEVEL

(exact order of this whole section still has to be fixed together with Vera Hansper, Suggestion: similar structure as the ops manuals, describing not only the technical but also the operational set up allowing researchers to enter European collaborations. In order for seamless interoperation it is extremely important to have OLAs and a high degree of communication between the operations teams)

Xxxxxx

(check overlap with MS 405 and 408!)

(Quote Steven Newhouse:)

We need to demonstrate that the new resources do not compromise the reputation we have for a production infrastructure.

4.1. COMMUNICATION CHANNELS

Xxxx

4.2. OPERATION LEVEL AGREEMENTS (OLA)

Xxxxx (availability, reliability here)

4.3. PROCEDURE FOR ADDING A NEW NGI

Xxxxx (enhance for broader scope)

4.4. SECURITY

Xxxxx
See operational security procedure Milestone 405! [R1]

(Input so far from Guiseppe Misurelli)

4.4.1. The integration into the EGI-CSIRT group

(a large majority of people involved come from the ex EGEE ROC security contact) of security experts from other MW stacks (I'm thinking about security officers rather then developers)

2. EGI-CSIRT has it's own plan for the development of security tools and what has been developed so far is mostly based on the gLite stuff so a discussion on how to cover other MWs is needed as well (quite the same for the extension of the Nagios tool to other MW stacks)

Anyway, I'll be very happy to help you reviewing the document giving the NGI operational center viewpoint and helping Elisa as soon as I'll be back on August 23.

(need: contact EGI-CSIRT)