



EGI-InSPIRE

INTEGRATING RESOURCES INTO THE EGI PRODUCTION INFRASTRUCTURE

EU DELIVERABLE: MS407

Document identifier:	EGI-MS407-V0-8.odt
Date:	18/10/2010
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.

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, 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.

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
0.5	20/08/2010	First internal reviews	Michaela Lechner / KTH Tiziana Ferrari / EGI.eu Mario David / LIP
0.6	27/08/2010	Input on Nagios and many comments Functionality descriptions	Mathilde Romberg / FZJ Michaela Lechner / KTH Emir Imamagic
0.7	18/10/2010	Input on GGUS and support functionality description	Torsten Antoni / KIT Michaela Lechner / KTH
0.8	18/10/2020	Input on Accounting	John Gordin / RAL Michaela Lechner / KTH
1.0		First complete 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.....	8
1.1. PURPOSE.....	8
1.2. APPLICATION AREA.....	8
1.3. REFERENCES.....	8
1.4. DOCUMENT AMENDMENT PROCEDURE.....	10
1.5. TERMINOLOGY.....	10
2. EXECUTIVE SUMMARY.....	11
3. INTEGRATION AT OPERATIONAL TOOL LEVEL.....	11
3.1. OVERVIEW INTEROPERATION STATUS FOR OPERATIONAL TOOLS - MW.....	12
3.2. INTEROPERATION AT INFRASTRUCTURE LEVEL.....	12
3.3. DEFINITION AND DESCRIPTION OF A MANAGEMENT INTERFACE.....	13
3.3.1. <i>Management Interface Functionality</i>	13
3.3.2. <i>Management Interface Requirements</i>	13
3.3.3. <i>Integration of new Resources into GOCDB</i>	14
3.3.3.1. Integration of new MW service types.....	14
3.3.3.2. Integration of new non-MW resources types.....	15
3.3.3.3. Declaration of new resources of an already available resource type in GOCDB.....	15
3.3.3.4. Regular review of the list of available service types.....	15
3.3.3.5. Summary of the complete procedure.....	15
3.3.3.6. Integrating gLite resources in GOCDB.....	15
3.3.3.7. Integrating ARC resources in GOCDB.....	16
3.3.3.8. Integrating UNICORE resources in GOCDB.....	16
3.3.3.9. Integrating Globus resources in GOCDB.....	17
3.4. DEFINITION AND DESCRIPTION OF A MONITORING INTERFACE.....	18
3.4.1. <i>Monitoring Interface Functionality</i>	18
3.4.2. <i>Monitoring Interface Requirements</i>	18
3.4.3. <i>Interoperability of different MW stacks with Nagios</i>	19
3.4.3.1. Critical tests and Nagios probes for gLite resources.....	20
3.4.3.2. Critical tests and Nagios probes for ARC resources.....	20
3.4.3.3. Critical tests and Nagios probes for UNICORE resources.....	20
3.4.3.4. Critical tests and Nagios probes for Globus resources.....	21
3.5. DEFINITION AND DESCRIPTION OF AN ACCOUNTING INTERFACE.....	21
3.5.1. <i>Accounting Interface Functionalities</i>	21
3.5.2. <i>Accounting Interface Requirements</i>	21
3.5.3. <i>Accounting</i>	21
3.5.3.1. gLite resources.....	22
3.5.3.2. ARC resources.....	22
3.5.3.3. UNICORE resources.....	22
3.5.3.4. Globus resources.....	22
3.6. DEFINITION AND DESCRIPTION OF A SUPPORT INTERFACE.....	22
3.6.1. <i>Support Interface Functionality</i>	22
3.6.2. <i>Support Interface Requirements</i>	23
3.6.3. <i>Integration of new resources into GGUS</i>	24
3.6.3.1. Integrating a new resource centre into the infrastructure.....	24
3.6.3.2. Integrating a new NGI in into the infrastructure.....	24
3.6.3.3. Integration of a new technology provider into the support infrastructure.....	24
3.7. DEFINITION AND DESCRIPTION OF A DASHBOARD INTERFACE.....	25
3.7.1. <i>Dashboard Interface Functionality</i>	25
3.7.2. <i>Dashboard Interface Requirements</i>	25
3.7.3. <i>Operational Dashboard Portal</i>	25
3.7.3.1. Integration of a new resource.....	26
3.7.3.2. gLite resources in the Operational Dashboard.....	27
3.7.3.3. ARC resources in the Operational Dashboard.....	27
3.7.3.4. UNICORE resources in the Operational Dashboard.....	27
3.7.3.5. Globus resources in the Operational Dashboard.....	27

3.8. USER MANAGEMENT, AUTHENTICATION AND AUTHORISATION.....	27
3.8.1. <i>User management in gLite and ARC</i>	28
3.8.2. <i>User management in UNICORE</i>	28
3.8.3. <i>User management in Globus</i>	29
3.9. INTEROPERATION BETWEEN OPERATIONAL STACKS.....	29
3.9.1. <i>Job Submission</i>	29
3.9.1.1. Direct submission from gLite-WMS to an ARC-CE.....	29
3.9.1.2. Direct submission from ARC to a gLite-CREAM-CE.....	29
3.9.2. <i>Data Management and Storage Infrastructure</i>	29
3.9.3. <i>Logging</i>	29
3.9.3.1. Real Time Monitoring of ARC resources.....	30
3.10. REQUIREMENT LISTS TO THE MIDDLEWARE PROVIDERS.....	30
3.10.1. <i>gLite</i>	30
3.10.2. <i>ARC</i>	30
3.10.3. <i>UNICORE</i>	30
3.10.4. <i>Globus</i>	30
4. INTEROPERATION AT PROCEDURES AND POLICY LEVEL.....	30
4.1. COMMUNICATION CHANNELS.....	30
4.2. OPERATION LEVEL AGREEMENTS (OLA).....	30
4.3. PROCEDURE FOR ADDING A NEW NGI.....	31
4.4. SECURITY.....	31
4.4.1. <i>The integration into the EGI-CSIRT group</i>	31
5. OUTLOOK AND FUTURE PLANS.....	31

Index of Tables

Table 1: Table of references.....8
Table 2: Glossary of terms.....10
Table 3: Outlining the current status of interoperation for each MW stack relative to the current set of operational tools.....12

Index of Figures

Fig. 1: Global architecture of the Operations Portal.....26

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 and described in their basic functionality.

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 what is considered to be basic operational tools is inherited from the EGEE project series experiences. However this might change in the future. Still we take this as a starting point when comparing the interoperability of different middleware components for each operational tool in our current horizon.

Operational procedures and policies are needed as well to enforce the application of the agreed basic set of operational interfaces to be supported by all resources. Some of the old EGEEIII procedures and policies may be adapted to the EGI era, while new requirements will have to be identified and turned into new procedures and policies. 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: https://wiki.egi.eu/wiki/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://infforge.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
R 32	UNICORE architecture http://www.unicore.eu/unicore/architecture.php
R 33	Relational Grid Monitoring Architecture http://www.r-gma.org/
R 34	APEL Home http://goc.grid.sinica.edu.tw/gocwiki/ApelHome
R 35	http://forge.ggf.org/sf/docman/do/listDocuments/projects.ur-wg/docman.root.current_drafts.aggregate_ur_schema
R 36	GGUS Documentation https://gus.fzk.de/pages/docu.php
R 37	NGI Creation Process https://wiki.egi.eu/wiki/Operations:NewNGIs_creation

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
GOCDDB	Grid Operations Centre DataBase
MPI	Message Passing Interface
OSG	Open Science Grid
R-GMA	Relational Grid Monitoring Architecture
RUS	Resource Usage Service
UNICORE	Uniform Interface for Computing Resources (middleware stack)
XUADB	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 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 currently existing operational tools in our scope. The actual operational tools might change though in the future.

Based on this table we gather 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 tools features to be fully integrated in the EGI infrastructure. Requirements can also stem from a more general interoperability point of view.

For completeness we also include a section describing the level of interoperation level between the different MW stacks, since this will get more important in the future when sites want to support several operational/middleware stacks in parallel, or for the interoperability of sites running different MW stacks.

Additionally, this document will give an overview of the status of operational procedures and policies needed for the integration of new resources and conclude with some future plans.

3. INTEGRATION AT OPERATIONAL TOOL LEVEL

EGI-InSPIRE scans eagerly the blueprint on how to successfully run a common European Grid infrastructure as inherited by the EGEE project series. A certain amount of optimization as well as decisiveness to pick up 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 nationally and regionally 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 deployed several middleware stacks. 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
GOCDDB	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	completed		not (yet) available	not (yet) available
3 rd level support in GGUS	completed			to be done

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

3.2. INTEROPERATION AT INFRASTRUCTURE LEVEL

The basic operational interfaces that must be supported for resources to be integrated into the EGI production infrastructure consist of a management interface, a monitoring interface, an accounting interface, a support interface and an additional graphical dashboard interface which collects and presents the information provided by the others and ties them together in a meaningful way to facilitate daily oversight grid monitoring duties.

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 GOCDDB is the tool of choice for fulfilling these management 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 new types of services provided by these new resources in GOCDDB.

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 have to be defined for these new services. 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 GOCDDB. 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.

3.3. DEFINITION AND DESCRIPTION OF A MANAGEMENT INTERFACE

3.3.1. Management Interface Functionality

A management interface is a kind of operational interface which allows to store, maintain and view the topology of the whole EGI infrastructure and the basic information of its resources.

Such an EGI management interface contains information and a placement in the topology order on:

- Participating National Grid Initiatives (NGI) and possible other groups (Countries, ROCs) and related information
- Grid Sites providing resources to the infrastructure including management, technical and security related contact points
- Resources and services, including maintenance plans and service status information access points for these resources
- Participating people, and their roles within EGI operations

Besides providing a central management tool to view and define production state, downtimes and maintenance status and whether a resource needs monitoring, it shall in essence depict what services are running where and who to contact for certain type of issues. The presented information can be a combined view of different regionalized or otherwise separated instances with their own local inputs.

3.3.2. Management Interface Requirements

The desired EGI management interface is requested to fulfil the functionality as described above. System and security contacts and higher level organisational management contacts for a site shall be easily identified. The management interface may provide a finer granularity in the people details by marking extended expertise on a specific middleware stack or an affinity to a certain type of service(s).

Additionally, it must be possible to register new kind of service types, groups or sites within the management interface.

We expect to have a role based interaction model to such a database, so that people responsible for certain sites, services or resources can update and maintain the various entries representing the entities under their responsibility within typical daily operations scenarios. In particular, basic service status information shall be easily viewable and changeable. It shall be easily possible to register a service of a known service type, to edit system administration information and put whole sites or single resources in and out of downtime according to predefined procedures. It shall be easy to identify whether a resource is monitored or not by the corresponding monitoring system. This monitoring bit can be set separately or implicitly within the number of production states. Decision on that has to be taken with hindsight to the fulfilment of practical use cases.

A management interface shall enable us to follow a resource through the certification process. The history and details of the certification process and other state transfers like uncification and suspension are desirable additional information.

Furthermore we expect a plug-in to an approved operational portal interface to be in existence or easily implementable due to using canonical standards.

Even though the information is mostly of static kind, a regionalized version with a central collecting portal of the management interface would of course be preferred in order to emphasize the distributed nature of the Grid community and to avoid single points of failure.

We follow up with GOCDB as a working example for an implementation of a management interface.

3.3.3. 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.

3.3.3.1. 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 3.3.3.4 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 3.3.3.3.

3.3.3.2. 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 3.3.3.1.

3.3.3.3. 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].

3.3.3.4. 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 unused 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 the Software Providers (EMI and EGI-JRA1) before producing a list of service types that are candidates for decommission.

3.3.3.5. 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.

3.3.3.6. 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...

)

3.3.3.7. 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. 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.

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

3.3.3.8. 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 (Sits in front of one or more UNICORE services as a gateway to the internet. Normally one Gateway per site.)
- Registry (All UNICORE services register here; clients ask the registry for available services in the Grid. Normally there's one Registry per Grid infrastructure. Backups can occur. The Registry works like a phone book and collects URLs of services.)
- Workflow Engine (Needed to add workflow functionality to UNICORE. Not needed if only single jobs are submitted within a Grid infrastructure. Normally there's one Workflow engine per Grid infrastructure.)
- Service Orchestrator (Handles dispatching of a workflow's subjobs, and brokering. One Service Orchestrator per Grid infrastructure)
- UNICORE/X (Hosts the XNJS, which handles job submission, file transfer, job monitoring etc., and the CIP. One UNICORE/X per supercomputer/cluster.)
- CIS (Information service. Standalone service which collects information from the UNICORE/X. One per grid.)
- XUADB (User database. Maps certificates or DNs to user logins, roles etc. Services like the Workflow Engine and the UNICORE/X query the XUADB for authorisation.

Pretty flexible how many there are per Grid; each site running their own XUADB seems to be the most common setting.)

- UVOS - Serves the same function as XUADB but is much more advanced and flexible by supporting arbitrary attributes, groups, advanced authorization, and more. Usually one per grid, but may be replicated.
- Target System Interface (TSI) (The actual interface to the local batch system; submits jobs and goes with the UNICORE/X.)
- SIMON (standalone service which monitors UNICORE sites, mainly by periodically sending test jobs.)

Some of these services are quite tightly coupled, 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. A more detailed view on UNICORE architecture can be found in [R 32].

3.3.3.9. 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. storage endpoint and data transfer service for the Globus middleware stack (GridFTP).
3. certificate based interactive login service (gssshd).

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

3.4. DEFINITION AND DESCRIPTION OF A MONITORING INTERFACE

3.4.1. Monitoring Interface Functionality

A monitoring interface shall first of all monitor the resources within the EGI production infrastructure. Such a test execution and notification environment is therefore needed. Critical tests and execute probes.

Grid monitoring is needed to ensure the infrastructure's reliability and find causes of failures immediately. In the optimal case actual failure is avoided by fine tuning the tests so that warnings about soon needed maintenance can be send. Critical tests to monitor all mission-critical infrastructure components have to be defined and implemented as probes. In the event of failure notifications of the possible problem together with hints on how to solve the problem are sent to the technical staff and other relevant people allowing them to work on the problem before outages affect production and availability.

Alerts and warnings are delivered to IT staff via email and SMS. Multi-user notification escalation capabilities ensure alerts reach the attention of the right people.

The execution of probes can be rescheduled to test the solution of a problem.

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.

The design of the monitoring interface is scalable and a failover concept is in existence.

A good monitoring system monitors not only the network and the resources, but also the accessibility and functionality of the used operational tools.

3.4.2. Monitoring Interface Requirements

- Regionalization is an important factor since Grid in its nature is a distributed system. Monitoring should therefore be split into various instances running in each region and a central instance collecting results. From the technical perspective the distributed system avoids scalability issues as each instance covers much smaller amount of sites than a single central instance. From the operational perspective the NGI teams get much more control and responsibility over the whole monitoring process. Otherwise problems at a central location would reflect on the whole grid, any changes would require consultations with a central body under the control of a single team. If something goes wrong in a regionalized scenario it can be quickly solved locally without asking a central party to perform actions. Finally, a distributed system enables individual instances to tune the monitoring by introducing extended custom probes to monitor monitor custom services not covered by the generic profile. Also, individual instances can benefit from additional functionalities of the monitoring system such as direct email or text message notifications, extending monitoring on uncertified sites or direct scheduling of tests via web interface.
- Status and historical data should be accessible in a centralized portal. These historical records of outages, notifications, and alert response are relevant for later analysis.
- The monitoring interface shall provide availability calculation component.

- The generic probe profile which also works as a basis for availability calculation has to be checked in regular intervals on up-to-dateness. In particular, if the current set of probes fulfils all the needs or has to be extended or reorganized. New probes shall be identified and provided. This should happen in coordination with the software providers.
- Information shall be exchanged according to a given template or prevalent open standard. E.g. the use of ActiveMQ (or an alternative) transport protocol is recommended.
- It shall work as an input plugin for the Operational Portal.
- Additionally it would be desirable to add an additional level and to not only monitor the resources put also the availability of needed operational tools, like the different regional monitoring instances.

We continue with referring to the Service Availability Monitor (SAM) in the regionalized Nagios monitoring framework based version where each region runs its own instance. This Nagios monitoring framework based solution was redesigned and chosen in favour of the former centralised SAM submission framework by the WLCG Grid Service Monitoring working group which was deployed at CERN during the EGEE project series to monitor the infrastructure's resources before being decommissioned on June 23th 2010.

3.4.3. Interoperability of different MW stacks with Nagios

Nagios <http://www.nagios.org/documentation> is a well-known and mature monitoring system that enables organizations to identify and resolve IT infrastructure problems.

Out of the box Nagios can already monitor many different infrastructure components - including applications, services, operating systems, network protocols, system metrics, and network infrastructure. Furthermore, its extendable architecture allows easy integration with in-house and third-party applications. Hundreds of community-developed addons extend core functionality to ensure a faultless functioning of the entire infrastructure. New critical tests to monitor further mission-critical infrastructure components can be defined and deployed with freshly written probes for them.

Within the EGI production infrastructure the central instance of Nagios collects the results and provides a centralized MyEGI portal [[link here](#)] to access status and historical data.

The current design was finished within the operation Automation Team (OAT) group.

A special Nagios box was established at CERN with the purpose of monitoring ActiveMQ Brokers network and Nagios instances. CERN developed probes for monitoring these two services. CERN committed to run this instance during the EGI project. The ops-monitor Nagios instance can be found on the following address: <https://ops-monitor.cern.ch/nagios>. Other operational tools developers were requested to provide probes for monitoring their tools as well. Once the probes are provided, they will be integrated into the ops-monitor Nagios instance. Details can be found in the RT ticket: <https://rt.egi.eu/rt/Ticket/Display.html?id=79>.

Analysis of failover configuration of centralized tools was performed. SAM/Nagios instances are supposed to be deployed at each NGI. Each NGI is responsible for fault tolerance implementations. Certain procedures ensuring that a NGI's Nagios is not down for a longer period of time are still needed.

To integrate a new MW stack into Nagios, critical tests for the service types defined in the management interface for this MW have to be defined and then Nagios probes for them have to be written. Possibly it is also sufficient to just have a compatible Nagios reporter from a different kind of monitoring tool which can be integrated in regional and central instances.

3.4.3.1. Critical tests and Nagios probes for gLite resources

Currently the Nagios probes for the following gLite service types are implemented:

- BDII (top and site BDII)
- CE
- CREAM-CE
- FTS_oracle
- LB
- LFC_mysql/oracle
- MON
- PX
- SE_dcache
- SE_dpm_disk
- SE_dpm_mysql
- VOBOX
- VOMS

Regarding these probes, further documentation and descriptions are found on twiki: <https://twiki.cern.ch/twiki/bin/view/LCG/SAMProbesMetrics>

3.4.3.2. Critical tests and Nagios probes for ARC resources

Historically Nagios' predator the former Service Availability Monitoring framework, 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.

3.4.3.3. 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].

3.4.3.4. 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.

3.5. DEFINITION AND DESCRIPTION OF AN ACCOUNTING INTERFACE

3.5.1. Accounting Interface Functionalities

The EGI Accounting Infrastructure collects CPU accounting records from sites and/or grid infrastructures and summarises the data by site, date (especially by month), VO, and user. This summary data can be displayed in a dedicated Accounting Portal by dynamic queries on the parameters above at any level of the hierarchical tree structure which defines EGI and partner Grids.

Accounting is necessary to demonstrate that delivered computing resources to a specific project are in accordance with signed Memorandum of Understanding agreements. Site

administrators are able to check actual usage of CPU resources against scheduling policies implemented at the site. VO resource managers are able to understand how CPU resources are utilised by their users.

3.5.2. Accounting Interface Requirements

An accounting interface has to fulfil the functionality described above. Further requirements are:

- Users need to be authenticated and be identified through to be able to collect user-level accounting information they are allowed to have access to.
- 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. Regional versions of the accounting portals are therefore necessary.
- Usage Records (URs) should comply to a common standard usage record if possible.
- A common transport mechanism needs to be identified to transport records across sites deploying different middleware stacks.
- Accounting of MPI jobs as well as accounting of virtual resources (grid-cloud integration) should be possible.

3.5.3. Accounting

The EGI Accounting Infrastructure is based on APEL [R 34]. ~~The it collects CPU and accounting records and the from sites and/or grid infrastructures and summarises the data~~ summerised by site, date (~~month~~), VO, and ~~u~~User. ~~This summary data can be~~ displayed in the Accounting Portal (internal and external refs http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php ?) by dynamic queries on the parameters above at any level of the hierarchical tree structure which defines EGI and partner Grids.

The bulk of existing sites collect data from their batch systems (list those supported), join it with the job's user grid credentials and publish to the central APEL repository. At the time of writing the EGI infrastructure is in transition of transport layer from R-GMA [R 33] to ActiveMQ already used by other EGI Operational Tools. Other partner Grids (list), and a few sites, with their own accounting services, publish summaries of data in the form described above to the APEL central repository. Sub-Grids of EGI (e.g. [Italian Grid Infrastructure](#) IGI) publish all of their VOs data. Partner Grids (e.g. OSG) publish selective VOs. In particular the LHC VOs are all published to APEL so that there is a single worldwide repository for LHC. At the time of writing, summary publishing is done by remote database insertion but an ActiveMQ summary publisher is under development.

CPU data are published in the form of either: job level records containing data from a single batch job; or summary aggregate records containing totals for a number of jobs run at a single site for a single user and VO in a given month. The Job Record schema is a plain text version of the OGF-UR v1.0 with some common extensions (~~footnote, since~~ the original UR did not have the concept of a site, which is so crucial to the Grid). The summary record has been submitted to OGF's UR-WG for adoption[R 35].

In addition to the ActiveMQ route for receiving data, the APEL development roadmap plans to have a RUS (ref, footnote WS [J. Ainsworth, S. Newhouse, and J. MacLaren. Resource Usage Service \(RUS\) based](#)

on [WS-I Basic Profile 1.0. UR, August 2005. ?](#)) interface to receive data only.

3.5.4. Accounting interface to ~~III~~ How wother infrastructures interface?

Other grid infrastructures who wish to publish accounting data need to:

- a) Define a structure for their Grid in GOCDB (or equivalent) that can be used by the accounting portal to display the data. The minimum requirement is a flat set of site names, used in the accounting records. (e.g. for OSG ~~the~~ [ise](#) data are obtained from MyOSG)
- b) Extract data from their accounting system grouped data by site/VO/User/FQAN/month and create each group into a 'summary record' meeting the APEL definition. Experience shows that for accounting systems using the OGF-UR this is a simple transform.
- c) Register the publisher with APEL (host DN). The ~~Ape~~[PEL](#) Repository only accepts accounting records from registered sites. For APEL client sites this is defined by the glite-APEL service type in in GOCDB. An equivalent mechanism will be developed for summary publishing sites/grids.
- d) Publish the records into the EGI ActiveMQ Message Bus with a destination of (queue/topic tbc). ~~The~~ APEL repository will accept the records into a holding container from where they will be merged with the summaries from other Grids and the summary produced by APEL from the job records it has received. Currently the master summary is rebuilt from scratch several times per day. Each time it uses the last set of summaries received from each Grid.
- e) ~~That's All Folks!~~ From the master summary table, the data are [then](#) exported to CESGA where they can be viewed in the portal. ~~(ref)~~

3.5.4.1. Future Work

At the time of writing the ActiveMQ interface into APEL only accepts a single type of job record. The summary development mentioned above will include handling multiple types of record. As well as the summary record this will allow the repository easily to be extended to support other types of accounting, such as storage, as well as allowing evolution of the cpu UR.

The RUS interface planned will allow other Grid infrastructure to use a standard web services interface to publish records. This will replace item (d) in the list above.

3.5.4.2. ARC resources

Accounting integration was performed already during EGEE III. The aim 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 used 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.

3.5.4.3. 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. [\(should mention more about RUS here as well?\)](#)

3.5.4.4. 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).

3.6. DEFINITION AND DESCRIPTION OF A SUPPORT INTERFACE

3.6.1. Support Interface Functionality

The user support infrastructure in use within EGI is a distributed infrastructure consisting of various topical and regional helpdesk systems that are linked together through a central integration platform, the GGUS helpdesk. This central total enables formalised communication between all partners involved in user support by providing an interface to which all other tools can connect and thus enabling central tracking of a problem, independent of the origin of the problem and the tool in which the work on the problem is done.

The interlinking of all the ticket systems in place throughout the project enables a passing of trouble tickets from one system to the other in a way that is transparent to the user. It also enables the communication and ticket assignment between experts from different areas (e.g. middleware experts and application experts) while at the same time allowing them to work with the tools they are used to.

A standard has been defined for the interface between ticket systems and also a template for a ticket layout exists to ensure the quality of service.

These are documented in the GGUS documentation [R 36].

For the case of EGEE and now EGI an own functional body has been introduced to keep track of the ticket processing management (TPM). The TPM keeps a global overview of the state of all tickets and is responsible for the non automatically assigned of tickets, so that they get forwarded to the right persons and the right units. The TPM teams also keeps track of long-term trouble tickets and helps to solve them with their very good general grid knowledge. Equally to TPM another 1st line support chain exists: the VO support team has members who have specific VO knowledge. In this way, a problem submitted to GGUS can be quickly identified as either a grid problem or a VO specific problem and addressed to the appropriate second line specialized support units.

The second line support is formed by many support units. Each support unit is formed from members who are specialists in various areas of grid middleware, or ROC supporters for operations problems, or VO specific supporters. The membership of the support units is maintained on mailing lists. A single e-mail address is available through which users can request GGUS for help. E-mails sent to this address are automatically converted into tickets and treated by the system.

3.6.2. Support Interface Requirements

Regardless of the number of parties involved, the submitter of a trouble ticket should be able to transparently follow the chain of actions needed to solve the initial problem. This transparency together with the independence from the actual ticket system used by the experts from the different areas who get assigned to the ticket can be seen as the main

requirements that ensure that information flows between different parts of the EGI infrastructure.

This is especially important since the support interface is not only used for 3rd level support dedicated to the end user, but also for relevant parts of internal trouble ticket communication fulfilling standard operational, grid oversight and partially also development functionalities.

Other relevant requirements on the support interface is the existence of a functional body like the TPM as described above and the connection to a useful, searchable and well maintained knowledge base.

Other basic requirements can be expected from any more advanced support ticket system:

- Differentiating between real problem tickets and service requests
- Ability to mark a ticket as spam
- Mail notification when a ticket is assigned to a support unit or person possible
- Possibility to involve several experts at the same time
- Searching tickets via ticket ID as well as via parameters
- Automatic reminders
- Several tickets describing the same problem can be put into a master-slave relation.
- Other dependencies can be represented with child and parent relations.

3.6.3. Integration of new resources into GGUS

There are three distinct cases to be considered when integrating new resources into the EGI user support infrastructure:

3.6.3.1. Integrating a new resource centre into the infrastructure

In case a new resource centre is added to the EGI infrastructure this is resources centre is always part of an NGI. This means that NGI management has to make sure that all steps are taken that are needed. For the user support area this is a simple case as the information about resource centres is extracted from GOCDB. This means that no manual steps are needed to integrate a new resource centre in GGUS.

3.6.3.2. Integrating a new NGI in into the infrastructure

If a new NGI joins the EGI infrastructure it is required to provide a ticket system which is integrated with GGUS. This can be done in different ways, depending of the size and the maturity of the NGI.

- The simplest way, which might be suitable for small upstarting NGIs is to use GGUS directly. This has the limitation of just one support unit for the whole NGI. Tickets cannot be assigned to specialised groups or specific resource centres within the NGI. This further processing of the tickets is done independently from the EGI support infrastructure.
- The NGI can make use of xGUS a customisable slimmed-down regional instance of GGUS. xGUS is hosted and maintained by the GGUS team. Customisation can be done via an administrative web interface, which enables creating and managing support units and defining special workflows. xGUS comes with the interface to GGUS built in.
- The NGI can set up its own ticket system. In this case the NGI has to make sure that their ticket system fulfils the requirements of the interface definition to GGUS. The NGI ticket system needs to be interfaced to GGUS and the NGI

is responsible for maintaining this interface. This for example includes testing the interface after releases of the GGUS portal.

Details on the NGI creation process can be found on a dedicated page in the wiki [R 37].

3.6.3.3. Integration of a new technology provider into the support infrastructure

Should EGI decide to utilise software from a technology provider that has not so far involved with the project, an agreement has to be found with that technology provider on who to integrate its support infrastructure with the EGI one. This process has taken place with EMI and IGE. No general rule how this will be done can be given here, as this is highly dependent on the internal support structure of the respective technology provider. Nevertheless it is important that this is done in a way that enables EGI to have an overview of issues with the products provided by the technology provider and to gather statistics on the quality of the support given by the provider.

EMI is about to set up a structure within GGUS for its various services, including e.g. 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].

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.

3.7. DEFINITION AND DESCRIPTION OF A DASHBOARD INTERFACE

3.7.1. Dashboard Interface Functionality

In order to operate a distributed infrastructure, management and monitoring information has to be collected and presented to ease the work of the operators of the infrastructure. The dashboard interface combines and harmonizes different static and dynamic information and enables the operators to react on alarms, interact with the sites, and provide 1st line support, as well as to really operate the sites and to supervise the creation and and the work on problem tickets on a regional and central level.

The dashboard allows predefined communication templates and is adaptable to different operational roles (1st line support, regional, central). Sites in the scope can be regional, central or predefined out of a list and can be sorted and displayed after several severity criterions to give an impression of not only one service put over needed actions for a whole region or even the whole EGI infrastructure.

3.7.2. Dashboard Interface Requirements

A dashboard interface has to fulfil the functionality described above. Further requirements are:

- access to a harmonized information service for...,
- access to a harmonized user authentication service for ... etc.

3.7.3. 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. Interlacing 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 Lavoisier 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 in 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).

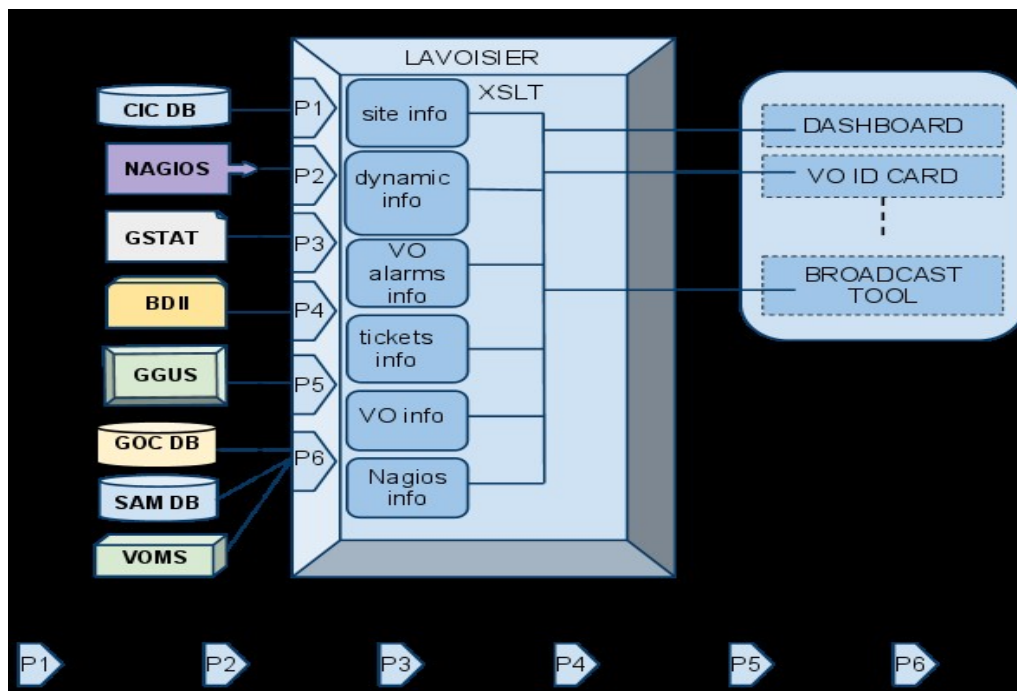


Fig. 1: Global architecture of the Operations Portal.

3.7.3.1. 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 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.

3.7.3.2. gLite resources in the Operational Dashboard

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

3.7.3.3. ARC resources in the Operational Dashboard

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

3.7.3.4. 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?)

3.7.3.5. 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.

3.8. 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)

(In EGI there are resource providers who will not offer pool accounts on their resources. Users have to apply for a personal account first and have a certificate mapped to it. To make life easier for the users within EGI a central service is needed where users apply for an EGI user account (within a VO) and then the accounts are created at the resource providers sites. Otherwise user would have to apply at each site for an user account and each site would have to generate the proper mappings.)

(maybe better as a sub point to accounting since we need to summarize data per user, might eventually be relevant to monitoring and operational portal as well though, but more obvious there?)

3.8.1. User management in gLite and ARC

VOMS is used for VO and user administration.

3.8.2. User management in UNICORE

The UNICORE User Database (XUADB) 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 XUADB user database. Every site can set up their own XUADB and an independent way of managing it, or there could be a central XUADB, or a central service that generates input for each site's XUADB (which is, so far, the most common way).

(Explanations and questions by Mathilde:)

The mapping of user DN (or certificate) to local userid at a single UNICORE site is provided by the XUADB.

The XUADB is a site local authorization component, maintained by the site.

In a Grid infrastructure using UNICORE with lots of sites there will be lots of XUUDBs. These XUUDBs have to be filled with the information of those users who are authorized to use the site's resource(s). How does a site get hold of this information in a large Grid infrastructure? Two examples: In D-Grid the central Grid Resource Registration Service (GRRS) knows about resources and which VOs are allowed to use them. Each VO has a VOMRS server where users are registered with their certificate and D-Grid Userid after they have applied for a userid and the VO membership. From these information a service is preparing mapping files for Globus, gLite, dCache, and UNICORE per site which then are used to feed e.g. the XUADB.

In DEISA users (who have been granted compute time on (a subset of) DEISA resources) apply at one of the DEISA sites for an account with their certificate. The DEISA user informations are collected in LDAP servers at the different sites that get synchronized once per day. Each site generates the input for its XUADB from its local LDAP server.

I'm sorry but I don't understand your XUADB comparison with GIIS as I only know that it stands for Grid Index Information Service, is a component of MDS, and maintains a current index of all the registered resources known in the grid. How does that relate to user entries in an LDAP with user information? The gstat command does what? I found the lcg qstat resources status display but isn't this site monitoring information and not user authorization information? Where is the link between your ldap (xuadb), giis and gstat? (the only thing I can imagine is that you are using an xuadb or an ldap service to authorize users to access certain parts of the information from gstat)

3.8.3. User management in Globus

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

3.9. 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.

3.9.1. 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.

3.9.1.1. 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.

3.9.1.2. 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].

3.9.2. Data Management and Storage Infrastructure

(dCache for ARC and gLite)

3.9.3. 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.

3.9.3.1. 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.

3.10. REQUIREMENT LISTS TO THE MIDDLEWARE PROVIDERS

Xxxxx (outgoing from points in 3.2, was thought as a kind of conclusion, we won't get any requirements from the NGIs yet.)

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

3.10.1. gLite

Xxxxx

3.10.2. ARC

Xxxxx

3.10.3. UNICORE

Xxxxx

3.10.4. 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 than 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)

5. OUTLOOK AND FUTURE PLANS

This milestone will improve over time. The functionality descriptions and the respective requirements of the different operational tool interfaces will definitely change over time and will have to be updated on a yearly basis.

One future plan action point is to collect operational requirements from NGIs that are interested in integrating novel resource types into their e-Infrastructure. Collection of requirements could happen

- during the EGI Technical Forum and/or
- through a set of dedicated face to face meetings with the most prominent NGIs in the fields in question.

Possible candidate list of relevant NGIs for such dedicated meetings are:

- Integration of UNICORE and Globus services: NGI-DE
Germany is the lead partner of the IGE project (<http://www.ige-project.eu/> about Globus support in Europe), the lead partner is Leibniz Supercomputing Centre
- Integration of desktop services: NGI-HU. Hungary is leading the EDGI project <http://edgi-project.eu/>
- Integration of cloud services: NGI-FRANCE and NGI-IBERGRID. France and Spain are involved in the StratusLab project
- Integration of storage resources into accounting: Italy and possibly other NGIs that are pioneers in this field. (Paolo Veronesi and Andrea Cristofori from the Italian Grid Infrastructure NGI), Paolo.Veronesi@cnaf.infn.it, Andrea.Cristofori@cnaf.infn.it)
- MPI accounting: Italy is certainly interested in this, together with Spain. Other NGIs from SEE region such as Turkey and Bulgaria have expertise/requirements in this, according to slide 8 of their presentation given at the SA1 kick-off meeting (<https://www.egi.eu/indico/sessionDisplay.py?sessionId=9&slotId=0&confId=43#2010-06-04>). The call for participation to the meeting will be open in such a way that any NGI that wishes to contribute is welcome.

We hope to present some of the gathered requirements already in the second version of this milestone. After collecting a set of requirements, the operational interfaces described in this milestone can evolve accordingly.

2) further action points for the coming months (I guess I would like to concentrate on procedures and basic standards)