



EGI-InSPIRE

OPERATIONS PORTAL WORK PLAN

EU MILESTONE: MS701

Document identifier:	EGI-MS701-final
Date:	05/11/2010
Activity:	JRA1
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/39

Abstract

This document defines the roadmap for the Operations Portal taking into account the regionalisation of the present CIC Operations Portal and the key operational tools, and new resource types being used on the infrastructure



I. COPYRIGHT NOTICE

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

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Cyril Lorphelin	CNRS	05/07/2010
Reviewed by	Moderator: Steven Newhouse Reviewers: Daniele Cesini	EGI.eu	06/08/2010
Approved by	AMB & PMB		01/09/2010

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1.0	05/07/2010	Draft	Cyril Lorphelin- CNRS
2.0	06/08/2010	Draft 2	Cyril Lorphelin- CNRS
3.0	01/09/2010	Final	Cyril Lorphelin- CNRS

IV. APPLICATION AREA

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



V. DOCUMENT AMENDMENT PROCEDURE

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



VI. TERMINOLOGY

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



VII. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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



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

VIII. EXECUTIVE SUMMARY

The Operations Portal is based on the “actor’s view” principle where each actor of the community has an access to information from an operational point of view according to his role in the project such as grid operator who daily monitors the status of resources and grid services, regular grid user and VO, site or NGI/ROC managers.

Complementary to this informative goal, the portal also fostered communication between different actors, through channels like broadcast notification, and downtime notification mechanisms and has set-up procedures to address their interaction needs.

It also offered an implementation of official operations procedures – like activity reporting for sites and ROCs or VO life cycle creation and update for VO Management. Last but not least, some integration of third party tools resulting from tight collaboration, like monitoring probes re-submission tool SAMAP (SAM Admin Page) or sites configuration tool (Yaim Configurator) proved useful to operators and sites.

The Operations Portal is currently divided into two instances:

- The historical Portal [R1]: the **CIC Operations Portal** that provides operational information related to VO, and deals with information related to sites and Regional Organization.
- The new portal [R2]: **Operations Portal** based on the framework Symphony [R3].

One of the applications in the historical portal, namely the dashboard for Operations, has been the first module to be migrated to the new portal. The migration is now complete and has included a complete re-engineering of the feature to cope with:

- tool regionalisation
- standard and generic developments (developments independent from the operating systems/environments and based on open source libraries and components)

The objectives for the coming year are the following:

- Migration of the remaining key features of the historical Operations Portal to Symphony (described in sections 2.1 , 2.2, 2.3)
- Provide, when possible, regional modules of these features (section 3)
- Provide a systematic standard access to all information handled by the Operations Portal Make the information available through standard format like xml for download or through API (section 5.3)
- Provide new adapters to Lavoisier [R4] Web Service to cope with new information sources (section 5.4)



- GOC DB harmonization (section 6)
- Provide further improvements to the existing features and customization as requested by the EGI Community.

The different timelines are summarized in the conclusion of this document (Table 1 - List of the timelines per features).



TABLE OF CONTENTS

1	INTRODUCTION	9
2	MIGRATION TO SYMFONY	11
2.1	VO Id Card.....	11
2.2	Broadcast Notification Tool	12
2.3	Downtime Notifications Mechanisms.....	13
3	REGIONAL PACKAGE	15
4	DASHBOARD	17
5	LAVOISIER 2.0	19
5.1	The role of Lavoisier	19
5.2	Re-factoring of the core engine.....	20
5.3	Programmatic Interface.....	21
5.4	Integration of new resources.....	22
6	GOC DB HARMONIZATION	24
7	ROADMAP SUMMARY.....	26
8	CONCLUSIONS.....	27
9	REFERENCES	28
10	APPENDIX : LIST OF FEATURES.....	30



1 INTRODUCTION

The historical Operations Portal has been built as an integration platform, allowing for strong interaction among existing tools with similar scope but also filling up gaps wherever functionality has been lacking. Moreover some features implement numerous work flows derived from procedures put in place during the EGEE project. These features came out of requirements expressed by end users or administrators of Virtual Organizations (VO), Regional Operations Centres (ROCs) or Resources Centres and Operations Coordination Centre (OCC) of EGEE.

The information on display is retrieved from several distributed static and dynamic sources – databases, Grid Information System, Web Services, etc. – and gathered within the portal. This has resulted in numerous tools that have become critical to sites like the User Tracking feature (to contact unknown users out of their DN in case of observed grid misuse for example) or the Alarm Notification feature (to subscribe to alerts upon monitoring failures). Linking this information in the portal has enabled us to display relevant and high added-value views of static and dynamic information of the production grid infrastructure.

The list of the different features implemented in the CIC portal and in the Operations Portal is available in Table 2 - List of the features with a short description and the current status.

This information can be provided by the central operational tools or by regional instances of these tools (See section 3). Authentication and authorization is based on user's roles associated to the user's certificate.

Complementary to the goal of securely processing information and workflows, the portal also fosters communication between different actors of the project, through channels such as the Broadcast and Downtime Notification mechanisms (described in section 4).

The architecture is the same for the both versions of the portal and is composed of three modules:

- A database – to store information related to the users or the VO
- A web module – graphical user interface – which is currently integrated into the Symfony framework (see section 2)
- A Data Aggregation and Unification Service named Lavoisier (see section 5)

EGEE Operations Portal purpose

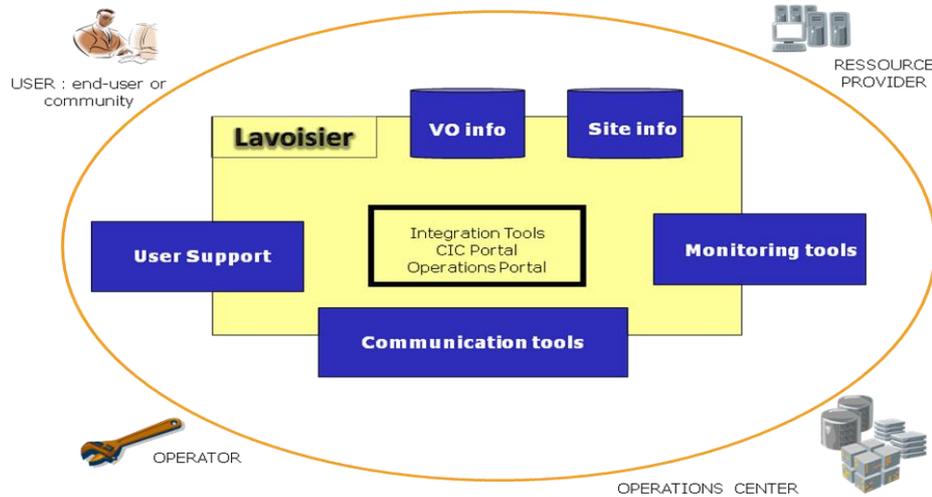


Figure 1- Operations portal an integration Platform based on Lavoisier

Lavoisier is the component used to store, consolidate and “feed” data into the web application. This module provides information from various sources without the portal being directly dependent on those information sources thanks to a caching mechanism. This indeed protects us from intermittent failures of information sources.

In **Figure 1**, we depict how the two instances of the Operations Portal (CIC Operations Portal and the new Operations Portal) are integration platforms for several heterogeneous information sources. These instances rely heavily upon Lavoisier to provide to users –being end-users, operators, resource centres or NGI/ROC.

2 MIGRATION TO SYMFONY

Symfony [R3] is a web application framework written in PHP which follows the model-view-controller (MVC) paradigm. Released under the MIT license, Symfony is free licensed software.

Symfony aims to speed up the creation and maintenance of web applications and to replace repetitive coding tasks. Symfony permits the development of a custom-made website built on proven techniques and standards.

Symfony is aimed at building robust applications in an enterprise context, and giving developers full control over the configuration: from the directory structure to the foreign libraries, almost everything can be customized. To match enterprise development guidelines, Symfony is bundled with additional tools to help developers test, debug and document projects.

We apply the framework approach for the new developments of the improvement of the features in the historical operations portal. The benefits we have seen until now are three-fold

- An increase in the efficiency of the Operations Portal application in term of response time to end-user requests.
- A decrease in the time spent on software maintenance.
- An increase in the robustness of the application by developing independent modules.

Our aim in the coming year is to continue migrating key features into this framework and to keep only one access point to information related to operations.

2.1 VO Id Card

The VO ID Card system records the life cycle of a given VO and links the VO managers to the project management for operations. The data is stored in the CIC DB, hosted at CC-IN2P3.

The schema of the database has evolved to support the VO life cycle and the evolution of the workflows supported by the portal [R5].

The VO ID cards are the static repository for VOs with information such as:

- The VO contact points (VO managers, VO User mailing list, VO representatives, etc.)
- The VO global information (enrolment URL, status, discipline, etc.)
- the Acceptable User Policy of the VO
- the VO's Core Services
- The VOMS information (Groups and roles, certificate details, etc.)
- Any other specified requirements (CPU needed, RAM needed, etc.)

The current interfaces permit a new VO to register and to subsequently update the information.



This registration and the updates are validated by the Chief Community Officer (CCO) of EGI.eu supported by the User Community Board. This replaces the Ressource Allocation Group (RAG) group used in the EGEE [R6]. Otherwise we use currently the registered roles (VO Managers and VO deputy) to restrict accesses to the VO ID card update.

The next Symphony ported version will include:

- a new “look and feel” of the features required in the EGI era
- the integration of new work-flows including
- yearly VO registration renewal
- workflow for VO deregistration

Moreover, we will continue providing enhancements of the current features:

- Distribution of available resources on a per VO basis through a synoptic display showing the geographical distribution of the resources for a given VO (e.g. CPU, disk/tape storage or other VO Services).
- all global information available per VO under several format HTML , RSS or XML
- statistics and metrics according to EGI’s evolving needs
- Notifications on demand about changes on the VO ID Cards.

The new work flow is still under discussion but developments have started, so that as soon as the work flow is validated the development can be finished within 3 weeks.

Estimated release date: August 2010.

2.2 Broadcast Notification Tool

With this tool every authenticated user (users with a grid certificates) are able to contact several categories of stakeholders impacted by a problem, an announcement, or a specific release. The aim is to provide to the different actors of the grid community with the information by mail or RSS feeds.

The current model is based on information coming from GOCDB [R7] (sites or ROC/NGI contacts) and from the CIC DB (VO contacts, mailing lists for operations).

In addition to sending out information we provide also:

- an archiving service
- a search engine by type of announcement – general or downtime notifications (section



13)- date, target e-mail, people contacted, body text, subject and author.

The next version in Symphony will provide:

- a new “look and feel” of the features revised for the EGI era
- the integration of new mailing lists provided by EGI
- the integration of templates
- the possibility to add customized contact lists
- the possibility to save personal settings for announcements for future use
- the possibility to attach a file to the broadcast.

Moreover we are studying the possibility to extend this tool with the user tracking module which can contact an end-user without having their email address, through his DN certificate. Indeed, we query the VOMS servers listed in the VO ID card, in agreement with the security policies, in order to collect all users email address.

Estimated release date: September 2010.

2.3 Downtime Notifications Mechanisms

The mechanism provided is based on the "supply and demand" rule. A site manager registers a downtime (supply) and grid users request to be notified or not (demand) via a subscription through the Operations Portal

When a site or region manager decides to register a downtime, it is done through the GOCDDB web interface. The relevant person fills an electronic form and specifies which site is concerned, the time lines for the downtime (start and end dates), and the reasons for the expected downtime. This information is stored in the GOCDDB databases. All downtime announcements are then collected and sent via the Operations Portal to all grid users who subscribed to receive downtime notifications according to his/her settings.

These subscriptions are done on a voluntary basis: every registered grid user is entitled to specify his/her downtime notifications settings via email, RSS feed or Ical feeds. A given user can create, modify and delete his subscriptions. When assigned a managerial role (i.e. ROC/NGI, RC, or VO manager) of a specific entity in the infrastructure (i.e. region/country, site, VO), a grid user can create, modify and delete subscriptions for the given entity in their scope.

A subscription also consists of one or more filtering rules. You can filter your subscription per scope: region, site, node, service, VO and all combinations of these parameters.



This solution is currently based on mail notifications or RSS feed. We want to extend this solution to technologies based on messaging system (Java Messaging System – ActiveMQ [R8]). Main advantage is to us to turn to a standard message broker within all our applications.

The main idea is to create on the fly the topics dedicated to a specific downtime notification at different levels:

- one topic about the downtimes declared for a NGI
- one topic about the downtimes declared for a site
- one topic about the downtimes declared for a service endpoint
- one topic about the downtimes impacting a specific VO.

A topic for each of the above or a combination of any of the above will be created as early as a downtime is scheduled and if this match the subscriber's pre-set preferences he will then be notified. Whereas at present, subscribers' preferences are triggering RSS feeds creation by default, regardless of existing notifications.

Estimated release date: December 2010.



3 REGIONAL PACKAGE

For around one year we have started the work on a regional package: the goal of this work is to distribute the different features and modules present into the Central Operations Portal into a regional package for NGIs to cope with NGI needs.

The first release of the first production package was on June 8th and this release included the first release of the dashboard module. The different part of the package and the documentation are distributed via a SVN repository [R9].

The application is composed of

- a web service named Lavoisier configured to handle Nagios notifications, store and provide data cache from GOC DB, GGUS and to generate metrics reports about the use of the dashboard (alarms raised, tickets handled ...)
- a PHP web application to provide a user interface.

The regional instance is linked with the central instance of Lavoisier; and creation, update, delete of records are synchronized so as not to disrupt global oversight operations. Synchronisation is achieved through REST and SOAP and records are synchronized every 10 minutes by using php scripts.

Any problem detected during the synchronisation is reported in a mail sent to webmasters.

The central and the regional instances have been built on the same model to behave in the same way and to be easily interoperable. A validation procedure [R10] has been established to validate the regional instances when a given region expresses this need. All regions or NGIs are able to opt for either the central or the regional instance. We will continue to offer as now regional views for the future features – when it makes sense - to the NGIs.

As you can see on the schema the architecture is exactly the same on a NGI and central instance. The distributed components are the same, just the configuration changes:

- with some piece of codes in the PHP application that permits to distinguish a regional instance from a central one
- On the Lavoisier side to filter information and to keep only the information related to the NGI.

This regional package is currently using the same central data sources as the central instance but this package is designed to integrate the regional version of GOCDB and regional helpdesks if they are deployed. This regional package will be constantly improved during the coming year:

- To add some of the other features that will be integrated step by step into the Central Instance.
These features will be added one by one into the central instance and if relevant to a regional instance we will provide an update to the package.
- To ease upgrade and set-up of the package in the regions

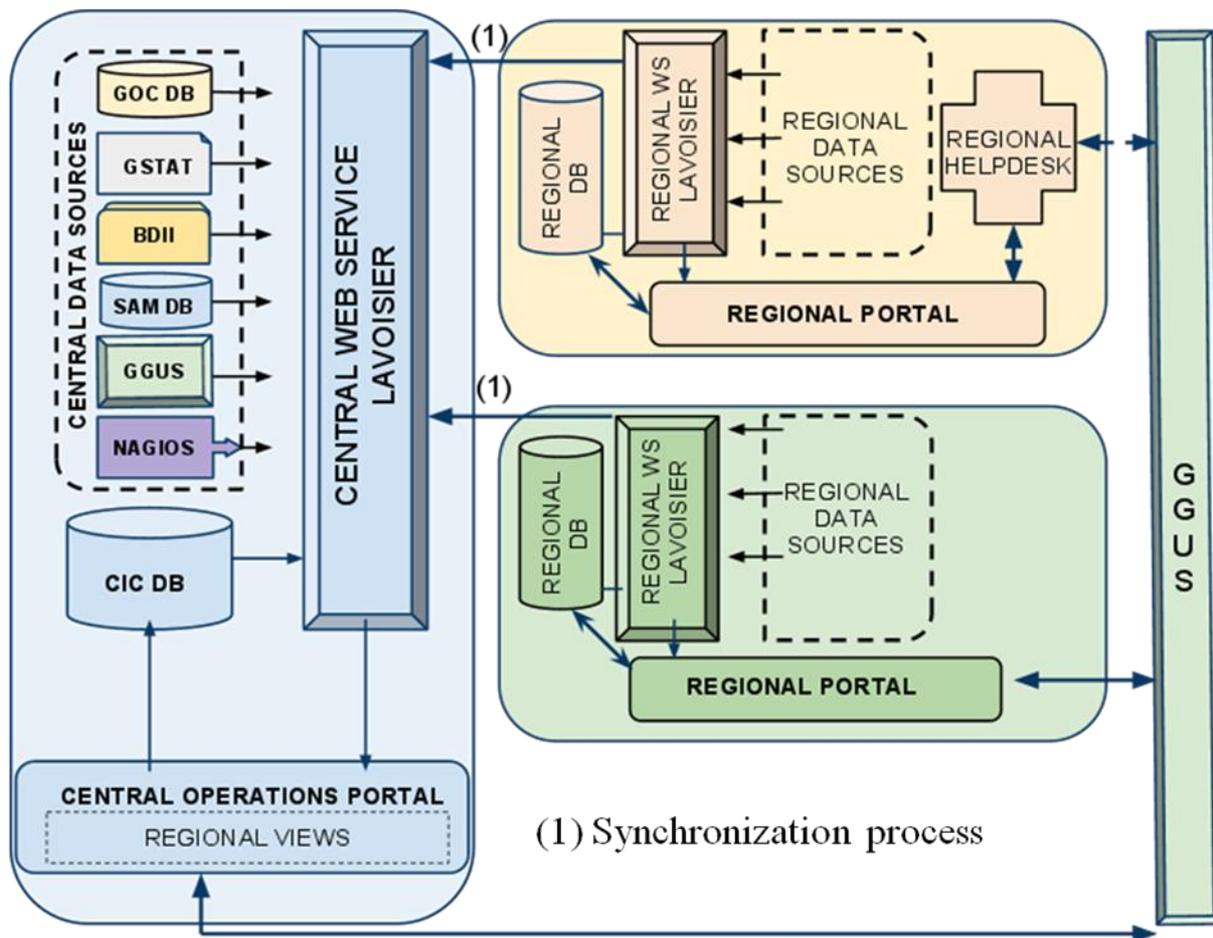


Figure 2- Interactions between the central and regional instances of the portal

For each central feature released, we will provide it in the regional package when applicable in the following months. We will also study the possibility as shown on the figure to use also regional helpdesks if they are based on the Web Service standards.



4 DASHBOARD

The dashboard is a tool designed to follow and track problems at sites.

This tool is an integration platform and provides a synoptic display of different data sources:

- Gstat [R11] , monitoring tool of the publication done by sites
- BDII , dynamic information published by sites
- SAM [R12], central framework of submission for VO specific tests

These results are only given only for VOs which have not migrated under the Nagios Tool. The VO specific tests are displayed as additional information in the dashboard

- Nagios [R13], the monitoring tool. This tool is the official reference used to monitor sites.
- GOCDDB , sites static database
- GGUS [R14], the EGI helpdesk:

Tickets created by Operators are part of the GGUS System. There is in GGUS a dedicated helpdesk for such tickets (CIC_Helpdesk). This is the one which is interfaced with the dashboard. This dedicated Helpdesk is connected with the global GGUS Helpdesk and tickets information transits from one to another. Having such a dedicated helpdesk ensures there is specific support for the needs of the operations staff (i.e. fields for site and node, specific escalation procedure).

In summary, operations staff can via a single dashboard interface can track problems using different results from the various Monitoring Tools (SAM , Nagios and Gstat) and can open or update trouble tickets.

We use also use the GOCDDB to consolidate monitoring information with downtime information, and GSTAT to provide dynamic statuses (Storage Usage, CPU Usage, number of jobs waiting and running). More information is found on the Dashboard how-to [R15].

As already mentioned, the dashboard has been the first module to be migrated under Symfony framework. The next steps of the dashboard evolution are as follows:

- Integration of monitoring results from other VOs, other than the OPS VO
 - through SAM VO specific tests – already in place for LHC VOs
 - through Nagios boxes dedicated to a VO

Some NGIs want to use other VOs than OPS to monitor their sites. For availability and reliability metrics a common VO is needed. So the solution is to set-up a separate Nagios Boxes with VO specific tests. We will replace step by step in the dashboard the monitoring of VO specific tests done by SAM by tests done by Nagios when they will be available.

Estimated Release Date: February 2011

Finally, we consider the integration of security monitoring via Pakiti [R16] results or dedicated security tests coming from the “security” Nagios Box. This is still under discussion at the time of



writing with the Security Monitoring Group [R17] as this group is currently working on the improvements of the Pakiti package and the development of new security probes for Nagios, which results should be exposed only to relevant people (i.e. grid security officers).



5 LAVOISIER 2.0

5.1 *The role of Lavoisier*

Lavoisier is an Open Source Tool developed at CC-IN2P3 which will be maintained in the long term. Lavoisier has proven effective in increasing the maintainability of the Operations Portal. Indeed, the unified view it provides has contributed to reduce considerably the amount and complexity of the portal's coding.

“Lavoisier” has been indeed 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 allows 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 maintainability is eased by making the code of the portal independent from the technologies used by the data sources and from the data cache management policy. Its design and interfaces made easier writing reusable code, and good performances are easily obtained by tuning the cache mechanisms in an absolutely transparent way from the portal code. Indeed, the different components work in a standardized way through the output of the Lavoisier Web Service. The translation of resource information into this standardized output is provided by different plug-ins.

Lavoisier is extensible; support for new data sources technologies can be easily integrated by developing new adaptors (or plug-ins). Some reusable adapters are provided to access data using various technologies, such as RDBMS, LDAP, Web Services, JMS, XML command line output stream, local and remote files (plain text, XML, HTML...), etc. Other reusable adapters take an existing data view and transform it to another data view, using technologies such as XSL, XQuery, SAX-based XML filtering. Introspection adapters expose data about data views configuration and current cache state. Efforts have been put into the configuration of Lavoisier, the structure of its caches and the rules of refreshing to have efficient, scalable and reliable data handling. Figure 3 describes the organization of the application around Lavoisier.

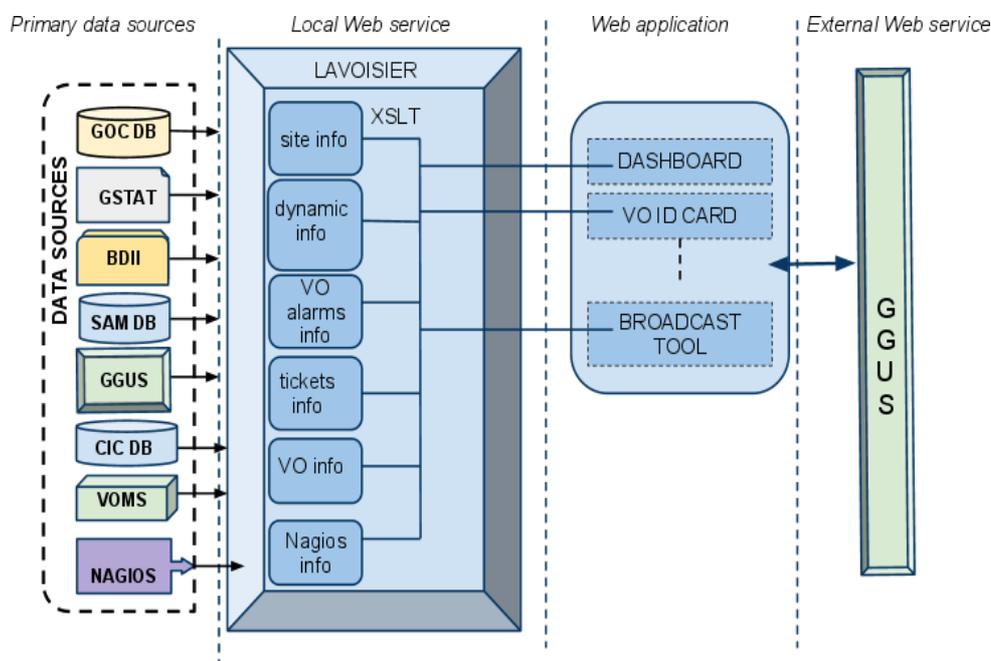


Figure 3 – Role of Lavoisier in the architecture of the application

The global information from the primary and heterogeneous data sources (e.g. GGUS, GOC DB, SAM, VOMS, Nagios) is retrieved with the use of the different plug-ins in Lavoisier. The information is structured and organized using the XSLT Language within Lavoisier. Finally, all of this information is available to the Web application, without the need for any further computations, which increases the efficiency of the Web Interfaces.

This architecture permits the web application to use the caches, even if a primary source is unavailable; hence we trigger access only to the information we need on the web page: The information is structured and therefore the primary data sources do need not be accessed hundreds of times but rather just a subset of the data that is already stored. Finally, we refresh the data sources only as needed and only when an action has been triggered. Last but not least, it is very easy to add a new data source in this model.

An effort will be put constantly to cope with this architecture to keep a flexible, modular and efficient application. The work of migration under Symfony modules is consequently complemented by a specific work on Lavoisier configuration to offer an easy and quick access to data.

5.2 Re-factoring of the core engine

Lavoisier has been increasingly used by the Operations Portal, and it is now reaching some limitations:

- Adaptors that consume a lot of memory when building the data view can overload the memory, because there is no mechanism to prevent several data views to be refreshed at the

same time nor mechanism to isolate a potentially problematic plug-in in a separate Java Virtual Machine.

- Since Lavoisier is more and more used intensively, it requires more and more optimisations (e.g. avoid useless serialisation/unserialisation of XML data, enable developing adaptors and features that consume few memory and CPU, etc.). Some optimized adaptors and features have been added progressively, but due to some Lavoisier limitations, they cannot always be used as easily and uniformly as other adaptors and features.
- Reusability of adaptors is limited by the plug-in interfaces and by the impossibility to chain adaptors that generate non-XML data views.
- The Lavoisier core engine cannot be easily extended. In order to improve flexibility and to benefit from development contributions, some currently hard-coded features such as cache content managers, cache refresh triggers and view validators needs to be implemented as plug-ins.
- Although configuration updates can be done on-the-fly, updating the code of an adaptor requires restarting the service.
- With the regionalisation of the CIC-Portal appears the need to improve communication capabilities between distributed Lavoisier instances, in particular notifications capabilities.

These problems drove us to perform a re-factoring of the core engine. We will introduce the use of a messaging system to manage events related to data views (access, cache refreshed, published...). The messaging system may be ActiveMQ, with in-memory messages for intra-Lavoisier communications, and OpenWire protocol for inter-Lavoisier communications.

This re-factoring is the pre-requisite to overcome the aforementioned limitations.
More about the upcoming features of Lavoisier 2.0 [R18].

Estimated Release Date: March 2011

5.3 Programmatic Interface

Lavoisier is currently used internally to “feed” the Operations Portal application. The next evolution is to provide uniform data access through Lavoisier to external consumer other applications than the Operations Portal. In order to do this without disturbing the Operations Portal application, we will set-up a second instance of the Lavoisier service by the end of the year. This instance will be usable through SOAP and REST protocols as the Programmatic Interface for other external applications than the Operations Portal. The aim of this Programmatic Interface is to act as a uniform data access provider.



Moreover the chosen standard is currently XML but we will provide also JSON on this programmatic interface in order to ease usage from Rich Internet Applications (AJAX, JavaScript).

Lavoisier currently provides the possibility to cross-query its data-views with XSLT language. We will provide XQuery and XPath as alternative languages. Both will be available in SOAP and REST operations. These additions will allow more flexibility on this programmatic interface.

Estimated Release Date: December 2010

5.4 Integration of new resources

The architecture of the portal has been built in order to extend the number of data sources, and to propose a standard access to information. This integration is eased by the use of Lavoisier.

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

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

Currently we have identified the needs for (1) optimizing existing plug-in and (2) to develop new plug-ins.

(1) A relational database management system plug-in is already in development. This plug-in will enable the user to see the entire database as a huge XML document. This XML document will not be built in order to avoid memory overload. Instead, the XPath query will be translated into SQL query, and only the result set will be converted to XML data. The structure of the document will be described by Object Relational Mapping files (Hibernate or Doctrine).

The use of this plug-in will enable cross-querying between the database and the other data sources, and could also be used easily via the programmatic interface by querying the data base through simple XPath queries.

(2) The integration of different Information Systems present into different middlewares such as ARC, UNICORE, or Globus will be done via an abstraction layer.

One possibility could be the SAGA Service Discovery specification [R19] (OGF) integrated into a Lavoisier plug-in which will permit to access information using different services (like the information service of UNICORE – CIS [R20]) and different schema like CIM [R21] or Glue Schema [R22] standards.

Lavoisier's flexibility allows us to be ready to integrate almost any kind of new information if needed and meaningful. For the new resource types coming into the EGI production infrastructure, such as HPC systems, virtualized resources, desktop resources, if these resources are monitored we are able to integrate its 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.

6 GOC DB HARMONIZATION

This work deals with the integration of GOCDB and the Operations Portal under a common interoperable toolkit for grid operations which will be done following two main directions:

- Integration of a common central human interface allowing users to access both central services through a single entry point.
- Integration of interoperable back ends for distribution to NGIs as a single package if possible.

Such development will require effort at data representation level as well as at interface and data transfer level.

Both teams aim at first at providing a common web interface to handle both type of information, being transparent to the user. This work is summarized by the transition between figure 4 and figure 5.

The solution envisioned has been RAL hosting of the GOC database and the Operations Portal in Lyon providing interfaces in the Symfony framework.

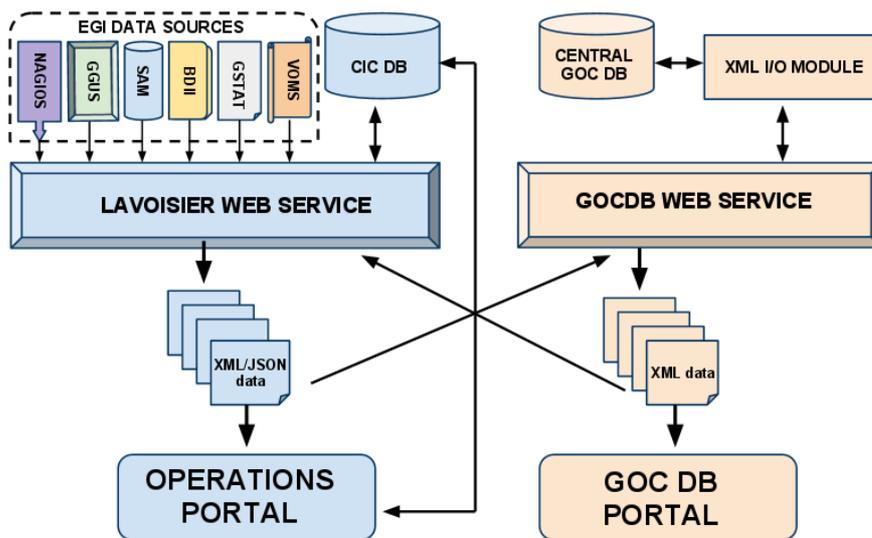


Figure 4 - Overview of the different components – Current status

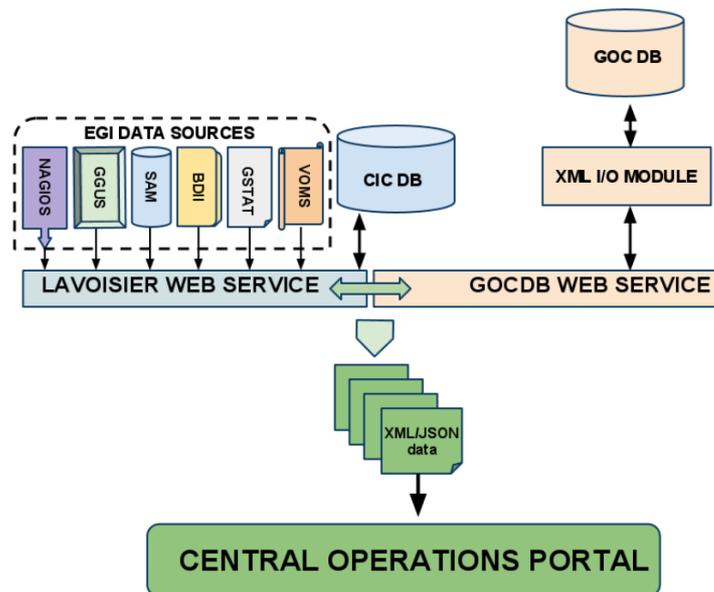


Figure 5 - Overview of the different components – Solution envisioned

Eventually this harmonization will be reflected in the regional package with a single portal, but also in a single package distributable in the different NGIs interfaced with the central “Operations Portal”. The complete merge of the instances (with the databases) is a complex process than needs to be studied.

A feasibility study of this integration has started and an assessment will be done shortly and the work will start in the first middle of next year. The initial focus will be on the central harmonization and the results will be propagated out it to the NGIs when they become available.

Estimated Release date:

Study of the integration: September 2010

Harmonization of central instances: August 2011

7 ROADMAP SUMMARY

This table is the summary of the roadmap of the main tasks considered for the coming year.

Features	Timeline
Migration to Symfony for the VO ID Card	August 2010
Migration to Symfony for the Broadcast	September 2010
Migration to Symfony for the Downtime notifications	December 2010
Enhancements of Lavoisier Programmatic Interface	December 2010
Integration of VO specific tests within the Dashboard	February 2011
Release of Lavoisier 2.0	March 2011
Study of the Harmonisation of GOC DB and Operations Portal	September 2010
Harmonisation of central instances of GOC DB and Operations Portal	August 2011

Table 1 - List of the timelines per features



8 CONCLUSIONS

The work of migration under a new framework is ongoing and we want to complete this part of the work as soon as possible. The main objective is to finish the migration by the end of the year (2010), so as to keep only one central point of access to information related to the EGI Operations.

This work should be strengthened with the integration of the GOC DB interfaces. In the meantime, the migration to Symfony provides the opportunity to adapt the different features to the EGI/NGI scopes.

This re-engineering will be focused also on a modular architecture based on the Lavoisier enhanced Web Service to ensure the flexibility and the extensibility of the application through the release of Lavoisier 2.0.

In parallel this component will allow us to deliver programmatic interfaces for external applications other than the Operations Portal and to step by step throughout the project integrate new types of resources coming from other DCIs (Distributed Computing Infrastructures).

9 REFERENCES

R 1	Operations Portal Historical Home Page : http://cic.gridops.org
R 2	Operations Portal New Home Page : https://operations-portal.in2p3.fr
R 3	Symfony home page: http://www.symfony-project.org/
R 4	Lavoisier Home page : http://grid.in2p3.fr/lavoisier
R 5	EGI VO Administration: https://edms.cern.ch/document/1070320/1
R 6	Resource Allocation Group wiki page: https://twiki.cern.ch/twiki/bin/view/EGEE/RAG
R 7	GOCDDB Home Page: https://goc.gridops.org
R 8	APACHE HomePage: http://activemq.apache.org/
R 9	SVN repository for Operations Portal Regional Package : https://cvs.in2p3.fr/operations-portal/package/
R 10	Validation of Operations Regional Package : https://cvs.in2p3.fr/operations-portal/package/validation.pdf?revision=HEAD
R 11	GSTAT Home Page: http://gstat-prod.cern.ch/gstat/about
R 12	SAM Programmatic Interface Documentation: https://twiki.cern.ch/twiki/bin/view/LCG/SamDbQuery
R 13	Nagios Home Page: http://www.nagios.org
R 14	GGUS Home Page: https://gus.fzk.de
R 15	Dashboard Documentation : https://edms.cern.ch/file/1015741/3/dashboardHOWTO.pdf
R 16	Pakiti home page: http://pakiti.sourceforge.net/
R 17	Security Monitoring Group: https://wiki.egi.eu/wiki/EGI_CSIRT:SMG
R 18	Lavoisier 2.0 : http://grid.in2p3.fr/software/lavoisier2/features.html



R 19	SAGA Service Discovery API: http://www.ggf.org/documents/GFD.144.pdf
R 20	Common Information Service : http://www.unicore.eu/community/development/CIS/cis.php
R 21	Common Information Model Home Page: http://www.dmtf.org/standards/cim/
R 22	Glue Schema specifications: www.ogf.org/documents/GFD.147.pdf

10 APPENDIX : LIST OF FEATURES

Features	Description	Status
Alert Notification	This feature enables to inform subscribers about monitoring alarms via an email at the same time these show up in the operations dashboard. A subscription page permits to refine the scope on which people are interested (site, service, VO).	This feature was based on SAM in the CIC portal and is now unavailable as SAM is decommissioned for OPS VO. A lot of feedback though has led us to provide this feature through RSS notifications based on Nagios records as a first step on the Operations Portal. A study will be done to put in place a notification system by email if needed.
Downtime overview	A synthetic view of the downtimes declared in GOC DB	This feature has been already integrated in the Operations Portal and graphics are generated dynamically.
Resources Distribution	Distribution of available resources on a per VO basis through a synoptic display showing the geographical distribution of the resources for a given VO , site or NGI	It should be migrated by the end of the year
Bazaar	This is an external application integrated in the CIC Portal. The Bazaar is a tool created to facilitate resource management in grid infrastructure. A main objective of the Bazaar is to enable better communication between two main actors in the process of resources allocation - resource providers and resource users.	The migration into the Operations Portal should be assessed with the tool developers and with the project
Broadcast	With this tool every authenticated user are able to contact several categories of stakeholders impacted by a problem, an announcement, or a specific release. The aim is to provide to the different actors of the grid community with the information by mail or RSS feeds.	Migration foreseen into the Operations Portal for September 2010

Features	Description	Status
Dashboard	With this application operations staff can via a single dashboard interface can track problems using different results from the various Monitoring Tools (SAM , Nagios and Gstat) and can open or update trouble tickets.	This feature has been already integrated in the Operations Portal
Documentation/ Procedures for Operations	The CIC portal is hosting or is displaying information about the operation activity.	Should be migrated to the EGI Wiki if the documentation is useful
Vo Management tool	The VO ID Card system records the life cycle of a given VO and links the VO managers to the project management for operations.	Migration foreseen into the Operations Portal for August 2010
ROC/RC reports	This feature was hosted until the end of EGEE in the CIC Portal. This tool was a reporting tool building a summary of the monitoring failures and alarms encountered in a specific period. Site/NGI managers were able to fill these reports for further problem analysis at later times.with their own explanations.	A migration of this in the Operations Portal has not been validated so far. Such an implementation could be considered if people feel the need for it.
Resources comparator	This feature enables to compare how a given resource is declared in GOC DB and/or BDII	The migration in the Operations Portal should be assessed.
Scheduled downtime Notification	A site manager registers a downtime and grid users request to be notified or not via a subscription through the Operations Portal	Migration foreseen for December 2010
User tracking	The user tracking tool is a module which can contact an end-user without having their email address, through his DN certificate. Indeed, we query the VOMS servers listed in the VO ID card, in agreement with the security policies, in order to collect all users email address.	The migration in the Operations Portal will be done by integration in the broadcast tool.

Features	Description	Status
YAIM VO configurator	This is an external application integrated in the CIC Portal. This tool lets you manage the last part of the YAIM configuration file of your site(s).	The migration into the Operations Portal should be assessed with the tool developers and with the project
Access to regional support unit	A summary of the regional support units via an interactive map	Not updated since the end of EGEE. The migration in the Operations Portal is not foreseen.

Table 2 - List of the features with a short description and the current status