

EGI grid middleware and distributed computing infrastructures integration

Emir Imamagic¹

University Computing Centre, University of Zagreb

J. Marohnica 5, Zagreb, Croatia

E-mail: emir.imamagic@srce.hr

Tiziana Ferrari

EGI.eu

Science Park 140, Amsterdam, The Netherlands

E-mail: tiziana.ferrari@egi.eu

The European Grid Infrastructure (EGI) is a distributed computing infrastructure (DCI) that provides more than 320,000 logical CPUs and 152 PB of disk space and data. EGI provides a set of basic operational interfaces that enable efficient management of the distributed computing environment and coordination of all stakeholders.

This paper describes the tools, interfaces and procedures that allow the integration of resources based on different software stacks into EGI. The integration relies on the operational interfaces of the tools that support daily operations: configuration management, monitoring, the operations dashboard, support and accounting. In parallel the operational interfaces are being extended in order to enable the integration with other resource infrastructures such as PRACE and EUDAT.

1

Speaker

The International Symposium on Grids and Clouds (ISGC) 2013
March 17-22, 2013
Academia Sinica, Taipei, Taiwan

1. Introduction

The European Grid Infrastructure (EGI) [8] is a publicly funded Distributed Computing Infrastructure (DCI) that provides more than 320,000 logical CPUs and 152 PB of disk space and data. EGI provides a set of basic operational interfaces that enable efficient and sustainable management of the distributed computing environment and coordination of all stakeholders (resource and infrastructure providers, technology providers, users). The basic operational interfaces consist of a management interface, a monitoring interface, an accounting interface, a support interface and a dashboard interface. These interfaces are implemented by tools and services that form part of the EGI operations portfolio.

The integration of resources whose access is supported by different software stacks or architectures is facilitated by a number of operational interfaces. Operational procedures and policies were developed to define the use of the agreed basic set of operational interfaces which should be supported by all resources. EGI provides access to resources using the following platforms: Advanced Resource Connector (ARC) [3], Desktop Grids [15], gLite [9], Globus Toolkit [14], QosCosGrid (QCG) [21] and UNICORE [25].

In the same time EGI operational interfaces are being extended in order to enable integration with other infrastructures such as PaRtnership for Advanced Computing in Europe (PRACE) and EUDAT, thus enabling users to easily utilize resources from multiple DCIs.

This paper describes the basic operational interfaces provided by EGI and the procedures for integrating new resources. Furthermore details are provided about the current integration status of the platforms supported by EGI and the work done to integrate with other infrastructures.

2. Operational tools

The EGI core infrastructure platform consists of set of operational tools which are used to assist all stakeholders (EGI operations, National Grid Infrastructure managers, site administrators, users, security staff, etc.) with their daily operations duties and to get and

provide support. These tools enable the smooth operations of a distributed federated infrastructure and operators' day-to-day activities.

The basic features provided by core infrastructure platform are: configuration management, monitoring, the dashboard, support and accounting. These features are provided by different tools which are developed and maintained by the EGI-InSPIRE project [8].

The following subsections provide a high level overview of each operational tool. Additional information can be found in [1] and [2].

2.1 Configuration management

The Grid Configuration Database (GOCDB) [5] contains general static information about the sites participating in one or more infrastructures. It provides topology information about those infrastructures by defining a respective list of service end-points. GOCDB allows sites to store, to update and view the topology of the production infrastructure and basic information about the respective resources within it. GOCDB acts as a reference database for all other operational tools, providing the definitive source of information about National Grid Infrastructures (NGI), sites, services and administrators responsible for services.

GOCDB can host information for multiple infrastructures and is independent of the services being registered. It can provide information about logical groups of services belonging to multiple sites. All data are stored in a database and users are provided with a common web interface. GOCDB is deployed as single, centrally managed repository within EGI, and is backed up to a remote failover instance every two hours.

Integration of new middleware stacks into GOCDB requires the creation of new service types. This step enables site administrators with resources running new middleware to register respective service endpoints. The appropriate procedure for adding new service type [6] is defined in order to unambiguously define the semantics associated to a service type name. New service types require documentation about the nature of the service requested. In addition the service type name reflects the usable scope of deployment (e.g. eu.egi.MPI, ch.cern.cvmfs.stratum.0).

2.2 Monitoring

The Service Availability Monitor (SAM) [20] is used to monitor the resources within the production infrastructure. SAM notifies site administrators in case of malfunctioning services and outages that concern their portion of the infrastructure. NGI operations teams and Virtual Organization (VO) operators are also notified for services in their scope. Monitoring results are also used for the calculation of availability and reliability of the infrastructure at different levels of aggregation. SAM is the authoritative source of alarms raised in the Dashboard of the EGI Operations portal.

SAM is based on well-established open source projects such as Nagios and the ActiveMQ messaging system. SAM consists of the following components:

- a probe execution framework based on the open source monitoring framework Nagios and the Nagios Configuration Generator (NCG),
- databases containing topology information gathered from GOCDB and other sources, test results and availability and reliability of sites and services,
- the Profile Management (POEM) system that provides an interface to group different tests into profiles which are the basis for configuring Nagios and other SAM components,
- the message bus infrastructure used for communication between distributed SAM instances,
- the visualization portal MyEGI that enables users to access current status, history and availability of monitored sites and services,
- a programmatic interface that enables other tools (e.g. Operations Portal, VO dashboards) to access test results and availability and reliability of sites and services,
- the probes used to test monitored services which are provided by middleware developers and third parties (e.g. NGIs, Nagios community).

SAM is deployed in a distributed manner. Each NGI deploys its own SAM instance that is responsible for monitoring sites and services of that NGI. The central instance (deployed at CERN) consumes results from all NGI instances, calculates availability and reliability and provides information for the central visualization portal. Figure 2 shows the architecture of SAM framework.

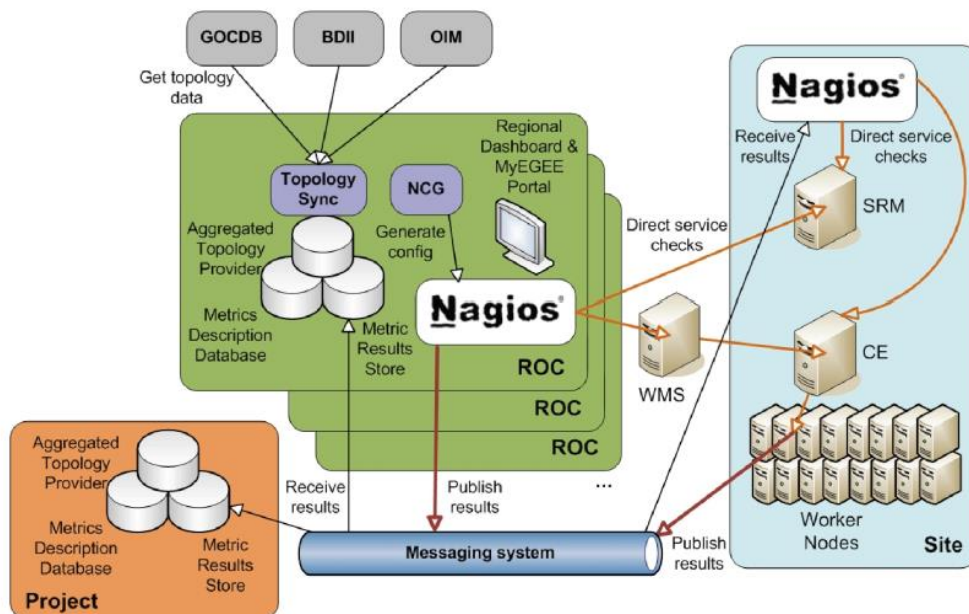


Figure 1 Architecture of the SAM framework

Integrating a new middleware stack into SAM requires the definition and implementation of tests for the service types defined in GOCDB. The probes are subsequently integrated into the SAM release. Integration also requires the probe developer to define the naming and test configuration (e.g. probe parameters, the execution frequency, timeout, etc). Two EGI procedures were defined to structure the integration activities:

- Adding new probes to SAM [11]
- Management of the EGI Operations Availability and Reliability Profile [12].

2.3 Dashboard

The Operations Portal combines and harmonizes different static and dynamic information and enables the operators to manage alarms coming from the SAM system. Operators use the dashboard to react on alarms, interact with sites, provide first-level support and perform oversight of alarms and manage tickets on national level.

The Operations Portal content is based on information which is retrieved from several different sources (e.g. GOCDB, the SAM monitoring system, EGI information discovery system, GGUS, web services, etc.) and made accessible through a single interface. The portal is organized around a highly modular web service Lavoisier that provides a transparent integration of various information sources. Figure 2 shows the architecture of the Operations Portal.

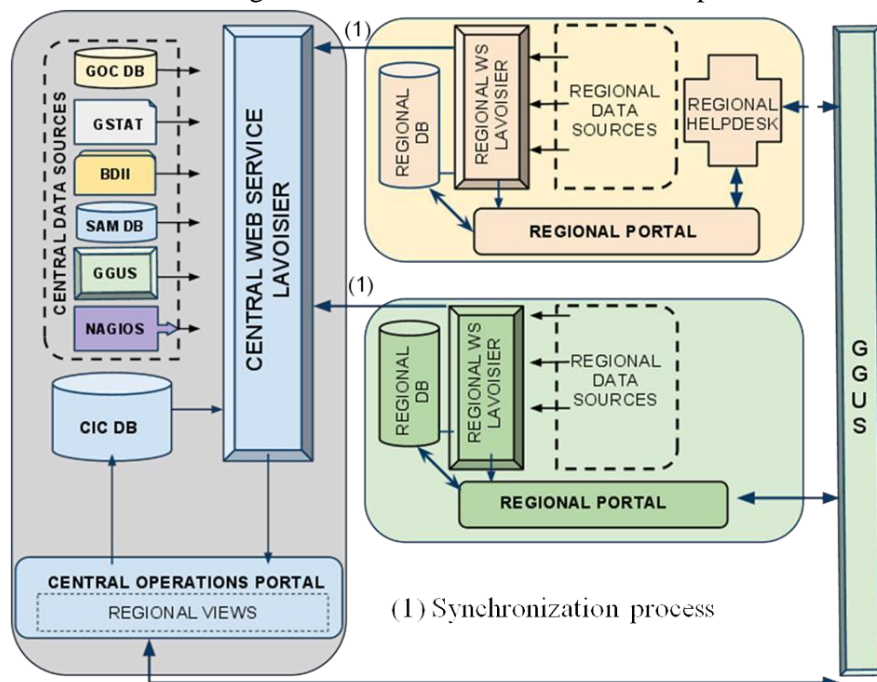


Figure 2 Architecture of the regional and central Operations Portal

The procedure “Setting a Nagios test status to operations” [13] defines how to add new test to a list tests which generate alarms in the Operations Portal. Once the test is added, alarms are

displayed in. These indicate problems with the respective service endpoints. Operators follow standard procedures for the handling of alarms.

If the resources using new grid middleware are already integrated in all the other operational tools through existing information providers (e.g. registered in GOCDB, monitored by SAM, present in the information system and registered in support interface), existing plugins can be reused and no additional integration effort is needed. For new technology providers, new plugins need to be developed and integrated.

2.4 Support

The user support infrastructure within EGI Helpdesk is distributed and consists of various topical and regional helpdesk systems that are linked together through a central integration platform, the Global Grid User Support (GGUS) system [7]. This central helpdesk enables formalized communication between the submitter of the incident record and all partners involved in user support by providing an interface to which all other tools can connect and enabling the central management of incidents.

The interlinking of all ticket systems in place throughout EGI enables to pass trouble tickets from one system to the other in a way that is transparent to the user. By exposing agreed interfaces, a hierarchy of helpdesk systems can be implemented allowing transparent exchange of incident records across different resource infrastructures. It also enables the communication and ticket assignment between experts from different areas (e.g. technology experts, operators and virtual research environment experts). A reference implementation was defined for the interface between ticket systems and also a template for a ticket layout exists to ensure the quality of service.

The regionalized implementation of GGUS is called xGUS. The xGUS is a simplified regional helpdesk instance that EGI can provide to resource infrastructure providers. These instances are operated centrally but can be customized by the regions by defining a specific support units. In xGUS the tickets can have a local or global scope. All answers to a global ticket are redirected to the central GGUS and centrally accessible.

The new middleware can use multiple approaches for integration into the EGI helpdesk. One approach is to instantiate a specific support unit in GGUS, which will enable operators to assign tickets directly to the appropriate support structure. This approach is heavyweight as middleware team must follow all EGI procedures. When new middleware is deployed in the infrastructure, an agreement has to be made with that technology provider on how to integrate its support infrastructure within the EGI Helpdesk.

2.5 Accounting

The EGI accounting infrastructure is a distributed system with central aggregation, where sites and infrastructures that are part of the federation publish their usage records centrally. The usage of the EGI computing resources by the users is accounted at service endpoint level and

the accounting data is stored in a central repository. In order to calculate the overall usage statistics of the distributed resource infrastructure, data from the different sites is aggregated. Usage records are transferred from services to the central repository by using the EGI message bus infrastructure.

The EGI accounting infrastructure currently collects CPU accounting records from sites and/or grid infrastructures and summarizes the data by site, date, VO, and user. Publication of accounting records for storage, virtual machines and parallel jobs is being validated. Summaries can be displayed in a central Accounting Portal, which enables a wide range of filtering options at any level of the hierarchical tree structure which defines EGI infrastructure (e.g. federation, national infrastructure and resource centre level).

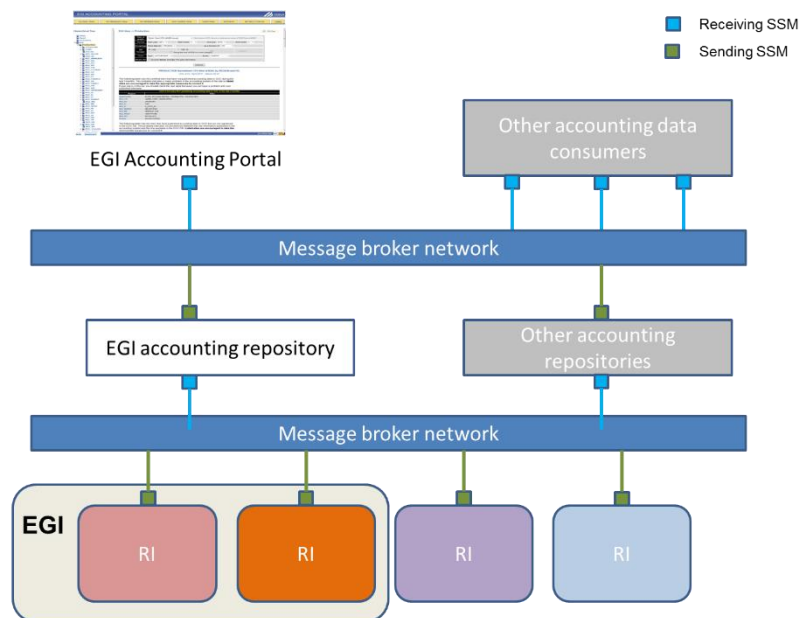


Figure 3 Exchange and collection of accounting records through SSM

The core EGI accounting infrastructure is based on APEL [4]. Other accounting systems have to interface to APEL to publish data into the central repository. Secure Stomp Messenger (SSM) [23] is the messaging system used by APEL to securely transmit messages. It is written in Python and uses the STOMP protocol [24]. Figure 3 shows diagram of data exchange with SSM. In addition, data must be formatted according to the OGF standard Usage Record scheme and an EGI-compatible profile. For a unified access to VO accounting information, sites can publish VO usage records centrally into the EGI accounting repository or into an alternative regional repository.

3. Software integration

The EGI infrastructure currently consists of resources that are accessible through the following middlewares: ARC, Desktop Grids, gLite, Globus Toolkit, QCG and UNICORE. Predecessors of the EGI infrastructure, DataGrid and later EGEE were heavily gLite-oriented, so gLite middleware is fully integrated with all operational tools. However, with the increased effort on integration other middleware are reaching a similar integration status.

The following subsections provide an overview of the integration status and the actions completed for each grid middleware. Additional information can be found in [2].

3.1 EMI middlewares

The European Middleware Initiative (EMI) [9] is a collaboration of the major European grid middleware providers: ARC, dCache, gLite and UNICORE. The main goal of the project is to improve the existing middleware and harmonize it, thus achieving a common framework with simpler and easier to use components. Such middleware at the same time reduces interoperability problems and simplifies deployment and maintenance for the distributed computing infrastructure providers. ARC and gLite are fully integrated into all EGI operational tools.

In order to speed up the integration of UNICORE resources, a task force within EGI was created. Thirteen service types for UNICORE services were added to the GOCDB. For monitoring purposes, Nagios probes were provided, which are maintained by the UNICORE project. In total there are ten tests fully integrated into SAM, out of which eight are integrated into the Operations Portal.

Accounting integration is possible with the publishers released with the EMI version 3.0 release which supports SSM version 2.0.

3.2 Globus Toolkit

The Globus Toolkit is an open source software toolkit used for building grids. It is being developed by the Globus Alliance and the project Initiative for Globus in Europe (IGE) [17]. IGE is especially focusing on providing packaged version of various Globus Toolkit components (e.g. MyProxy, GRAM, GridFTP).

In order to speed up the integration of Globus resources, a task force within EGI was created. Currently there are 5 Globus service types in the GOCDB. Standard Nagios probes and probes provided by the Croatian NGI are used for monitoring. Globus Toolkit resources are currently monitored with 8 different tests, all of which are integrated into the Operations Portal.

Globus Toolkit provides its own accounting system (GridSAFE). After the IGE 3.0 release integration of GridSAFE with APEL via SSM client was successfully tested.

3.3 Desktop Grids

The International Desktop Grid Federation (IDGF) [15] develops middleware that consolidates the results achieved in the EDGeS project concerning the extension of Service Grids with Desktop Grids (DGs) in order to support EGI and NGI user communities that are heavy users and require an extremely large number of CPUs and cores.

GOCDB was extended with three new services types representing desktop grid resources. A probe for monitoring Desktop Grids services was integrated into SAM. The new follow-up project International Desktop Grid Federation - Support Project (IDGF-SP) [16] is going to continue work on integrating into EGI dashboard, support and accounting interfaces.

3.4 QosCosGrid

The Multiscale APplications on EuRoPean e-infrastructures (MAPPER) project [19] aims to deploy a computational science environment for distributed multiscale computing, on and across European e-infrastructures. In order to further the project's aim, MAPPER initiated collaboration with EGI InSPIRE and PRACE. The main goal was to integrate two applications that perform distributed multiscale computing by using a set of core middleware services from the QCG middleware stack.

For the needs of QCG middleware 3 new service types were added to the GOCDB. QCG team provided set of probes for monitoring QCG services. Probes were integrated into SAM and 5 new tests were added both to SAM and to Operations Portal.

QCG team provided accounting solution compatible with APEL system. Deployment in production is possible with the publishers released with the EMI version 3.0 release which supports SSM version 2.0.

4. DCIs integration

EGI is currently investigating the operations integration with two different infrastructures: PRACE and EUDAT. PRACE [20] is a pan-European distributed supercomputing infrastructure, providing access to high performance computing resources. EUDAT [26] is a collaborative distributed infrastructure that enables users and communities to seamlessly share and process data.

Requirements were collected from MAPPER users which are in need of using both EGI and PRACE resources. The main requirements are:

- the need for a streamlined access to e-Infrastructure services and resources through common mechanisms for resource allocation reducing the bureaucratic overhead,
- the availability of monitoring information across EGI and PRACE,
- the availability of accounting information through a single portal,
- the user authentication through X.509 certificates,

- harmonized user support in case of problems concerning different infrastructures at the same time, i.e. single point of contact for both EGI and PRACE infrastructures.

A feasibility study of the integration of helpdesk and accounting infrastructures was identified as the first step of the EGI-PRACE collaboration. The user requirements mentioned above need common policies for secure access to confidential information. EGI and PRACE already collaborate on security matters by sharing security policies and communication channels to support integrated security incident response activities.

5. Conclusion

The EGI operations portfolio consists of services that are crucial for sustainable and efficient management of the EGI infrastructure. These operational services are designed in a modular fashion that enables integration of new grid and other softwares into EGI. Several procedures and policies are defined to ensure the smooth integration process. This paper provides overview of operational interfaces and current status of middlewares integration into EGI infrastructure.

Functionalities and requirements of the different operational interfaces will evolve over time, based on the input from resources providers and research infrastructures. Integration with other DCIs will also bring new requirements for the extension of the operational interfaces currently deployed in EGI.

References

- [1] S. Burke, T. Ferrari, M. Krakowian, P. Solagna, *EGI-InSPIRE D4.6 EGI Operations Infrastructure*, [<https://documents.egi.eu/document/1309>]
- [2] E. Imamagic, M. Lechner, *EGI-InSPIRE MS421 Integrating Resources into the EGI Production Infrastructure*, [<https://documents.egi.eu/document/1308>]
- [3] *Advanced Resource Connector (ARC)*, [<http://www.nordugrid.org/arc/>]
- [4] *APEL*, [<https://wiki.egi.eu/wiki/APEL>]
- [5] *Grid Configuration Database (GOCDB)*, [<https://wiki.egi.eu/wiki/GOCDB>]
- [6] *GOCDB Documentation – Adding new service types*, [https://wiki.egi.eu/wiki/GOCDB/Input_System_User_Documentation#Adding_new_services_types]
- [7] *Global Grid User Support (GGUS)*, [<https://wiki.egi.eu/wiki/GGUS>]
- [8] *European Grid Infrastructure (EGI)*, [<http://www.egi.eu>]
- [9] *European Middleware Initiative (EMI)*, [<http://www.eu-emi.eu>]

- [10] *EGI InSPIRE paper*, [<http://go.egi.eu/pdnon>]
- [11] *EGI Procedure – Adding new probes to SAM*, [<https://wiki.egi.eu/wiki/PROC07>]
- [12] *EGI Procedure – Management of the EGI OPS Availability and Reliability Profile*, [<https://wiki.egi.eu/wiki/PROC08>]
- [13] *EGI Procedure – Setting Nagios test status to operations*, [<https://wiki.egi.eu/wiki/PROC06>]
- [14] *Globus Toolkit*, [<http://www.globus.org>]
- [15] *IDGF*, [<http://www.desktopgridfederation.org>]
- [16] *IDGF-SP*, [<http://www.idgf-sp.eu>]
- [17] *Initiative for Globus in Europe (IGE)*, [<http://www.ige-project.eu>]
- [18] D. Lecarpentier, *Towards a collaborative data infrastructure for science*, ISGTW, 2012, [<http://www.isgtw.org/feature/towards-collaborative-data-infrastructure-science>]
- [19] *MAPPER*, [<http://www.mapper-project.eu/>]
- [20] *PRACE*, [<http://www.prace-ri.eu/>]
- [21] *QosCosGrid (QCG)*, [<http://www.qoscosgrid.org>]
- [22] *Service Availability Monitoring (SAM)*, [<https://wiki.egi.eu/wiki/SAM>]
- [23] *Secure Stomp Messenger (SSM)*, [<https://wiki.egi.eu/wiki/SSM>]
- [24] *Stomp Protocol Specification*, [<http://stomp.github.com/stomp-specification-1.1.html>]
- [25] *UNICORE*, [<http://www.unicore.eu/>]
- [26] *EUDAT*, [<http://www.eudat.eu/>]