e-infrastructure

PROPOSAL FOR A GRID OBSERVATORY SPECIALIZED SUPPORT CENTER – GOSSC

VERSION 2.1

Document identifier:	Document2
Date:	**16/07/2009**
Activity:	**EGI-applications**
Lead Partner:	**LRI-CNRS**
Document status:	**DRAFT**
Document link:	

Abstract: This document is a **preliminary** version of the GO-SSC activity for the EGI-applications proposal. The full work plans will continue to evolve towards a more precise specification with the partnership constitution and contributions.

Document Log

Issue	Date	Comment	Author/Partner
1.0	30/06/2009	Outline	LRI-CNRS
2.0	14/07/2009	First draft for review	LRI-Logica-MIS

TABLE OF CONTENTS

1. INTRODUCTION

1.1. PURPOSE

This document will present the overall proposal for the GO SSC, including motivation, strategic and governance aspects as well as detailed work plans.

1.2. DOCUMENT ORGANIZATION

The document follows the SSC Guideline V6 [R 1] . Note that the work plan and organization section are in a very preliminary status.

1.3. REFERENCES

R 1	SSC Guideline V6 https://edms.cern.ch/document/989620/1

Table 1. Table of references

2. EXECUTIVE SUMMARY

2.1. MOTIVATION

Scientific communities worldwide have set up massive grids that manage several tens of thousands of CPU's and several PetaBytes of storage space. The control, and maintenance of these complex systems remain a significant operational challenge. Application developers need synthetic characterizations of the grid activity and the grid applications for predicting and optimizing application performance. Grid models are required for dimensioning, capacity planning, or predicting the improvements consecutive to changes in grid configuration or middleware.

The grid infrastructure consists of a variety of hardware and software components, which are, in their own right, complex systems. Experimental data on the grid activity in real working conditions and advances in modelling method are necessary to discover adequate empirical models of the grid. On the other hand, fundamental grid research needs the experimental data created by the collective behaviour of the first grid users communities, as an input.

2.2. GOALS

The goal of the Grid Observatory SSC is to contribute to an ***experimental theory of large grid systems*** by setting up a ***data repository***, enabling the development of ***behavioural grid models*** and supporting the deployment of ***software experimental facilities***.

The Grid Observatory SSC includes the collection and publication of grid activity traces, the construction of an ontology of the domain knowledge, and the exploration of new grid models and control methodologies.

The GO SSC will build on the previous activity of the GO cluster in EGEE-III, which has already successfully realized a grid trace portal and fostered national and international collaborations with the Computer Science community.

Moreover, the GO SSC will act as a catalyst for developing synergies at the European level between scientific communities that have had so far limited opportunities to interact, with a special emphasis on the cross-fertilization of autonomic computing on one hand and grid research and engineering on the other hand.

The added value of the GO SSC is in the integration of its production goal – make available comprehensive and usable grid traces – and the long-term scientific goal of acquiring better knowledge and control of the grid as a complex system

An important asset of the GO is to provide the multidisciplinary skills and institutional positions required to make truly significant contributions towards this goal. The partnership includes 1) high-level scientific expertise in both Distributed Systems, Machine Learning, and Ontological Analysis 2) support from the EGI-related projects both in operations and in middleware development.

2.3. STRUCTURE

The GO SSC is intended to be a stable entity whose primary goal is safeguarding and publishing datasets in the long run. Consequently, it is extremely important that the GO SSC can evolve towards a permanent structure and define a sustainable financing model.

On the other hand, the activity is much younger than all other scientific SSC in the EGEE framework, and is still in it ramp-up phase; thus it requires initial development funding, and has to invent a permanent structure and a governance model in the course of its existence.

3. USER COMMUNITY

1.1 SCIENTIFIC AND ENGINEERING SCOPE

The scientific area covered by the SSC are the segments of Computer Science and engineering related to complex systems on one hand, and large scale distributed systems on the other hand. <2be developped>.

Industry is active in the area of Autonomic Computing. In order to ensure a production-quality result on the data collection and publication activity, as well as keeping pace with the developments in the cloud area, collaboration with industry and EU industrial projects has been seek. Discussion for partnership with IBM, Logica, SixSq, and the EU project RESERVOIR are ongoing.

1.2 GEOGRAPHIC SCOPE

Discussions for support on the current proposal are in progress with various NGIs. The topic has or will be discussed with the NGIs in France, Czezch republic, Italy, and UK, which are expected to be actively involved in the SSC.

The SSC aims at fostering international collaboration by enabling the research groups from different countries and regions to collaborate with the same data, to build common reference sets (benchmarks), and to propose operational challenges derived from real-world datasets.

1.3 AIMS

The aims of the SSC are as follows.

- Provide production-quality services for the Computer researchers and engineering in Europe and beyond, through data collection, publication and descriptive analysis.
- Foster basic research collaboration through scientific networking.
- Build bridges between the operational requirements emerging from the new EGI model of operations and the computer science community, encouraging and facilitating the experimentation of new innovative ideas contributing to grid middleware improvement in reliability, stability and performance.

More specific aims of the proposed GOSSC are listed below, with the related tasks in the work plan indicated

- Develop of a full-fledged acquisition process integrated within gLite (T1).
- Provide and develop on-line analysis services running on the EGI grid, as contributions to on one hand scientific research (T1), and to the other hand a sustainable operation model (T2).
- Provide a network of expertise in the interpretation of production grid behavioural models (T2).
- Contribute to the definition of a grid ontology, which will be the basis for interoperability with other data repositories (GWA,…) and interaction with other computational production models (clouds, grid overlay networks, desktop grids) (T3).
- Foster the creation of a COST project as a support for the basic research networking (T4).
- Define and enact processes for the specification of interpretation and control challenges , and the evaluation of the proposed solutions (T4)

1.4 LEGITIMACY

The SSC must have support from, on one hand stakeholders involved in actual production, such as some NGIs and EGI.eu, and on the other hand research institutes not presently involved in EGI, but prospective users of the GO data and services. The French NGI will provide the bulk hardware resource and participate in the maintenance tasks. The other NGIs are expected to contribute to the

acquisition and interpretation tasks, and thus to commit human resources in the SSC. The GO SSC requires interaction with the future EGI.eu, both at the operational level, in order to keep pace with the general development of the infrastructure, software monitoring resources, operational issues, and to evolve to a sustainable set of services.

Considering computer science users, the proposed SSC expects to receive Letters of Support from various high-level research groups, institutes and projects. As an example, discussions have started with the NSF Centre for Autonomic Computing, the French experimental infrastructure for Computer Science Aladdin/G5K, the Core Grid ERCIM Working group, and the RESERVOIR project. However, the research activity will be mainly funded by applying to relevant academic, regional, national and international calls, as has already been done.

4. GOVERNANCE MODEL

The GO SSC is intended to be a stable entity. Consequently, it is extremely important that the GO SSC can evolve towards a permanent structure. On the other hand,

- the activity is much younger than all other scientific SSC in the EGEE framework, and is still in it ramp-up phase;
- the Computer Science community has no international body comparable to CERN, ESA, or even the large biomedical collaborations.

The following tables represent the current status of the SSC partners and stakeholders. Italics are **prospective** partners/stakeholders. Table 2 lists the NGIs and the corresponding partners. Table 3 gives the list of industry partners. Table 4 gives European and international partnerships. **It must be stressed that the partnership is still in construction.**

NGIs	Funded partners	Networking and expertise[1]	Community
France	CNRS (LRI, MIS, LAL, LPC)		*Aladdin/G5K (Inria)*
UK	Imperial College	*U. Cardiff*	
Italy	UNIPM	*U. Bologna*	
Czech Republic		*Cesnet*	
?	Partners from Core Grid		

Table 2. Partners, NGI view

IBM
Logica
SixSq.
HealthGrid

Table 3. Industry partners

Project or Institution	Status
Center for autonomic computing (CAC)	NSF center
CoreGrid	*ERCIM Working Group*
Reservoir	*EU project*
EGI operations	planned EU project
EMI	planned EU project

Table 4. International partners

The usage of the SSC products is, by definition of the SSC, open to the whole scientific community. The governance model issue is thus limited to the steering of the SSC, both scientific and resource-wise.

- Resource-wise, the funded partners are the first concerned, and the rules of FP projects will apply as usual.
- Concerning the scientific steering, one of the goals of the SSC is precisely to catalyze the interaction between different communities. Thus defining a governance model will be possible only as a result, and not a prerequisite, of the SSC activity. As a first formalization, a MoU will be elaborated, with the French Institut des Grilles as the hosting institution. All permanent institutions (ie the listed NGIs, and CAC) will be solicited for endorsing this MoU. Moreover, the OGF will also be solicited.

How other EU projects can be formally involved, and their role in the governance model, is a transversal question to all SSCs. <2 be completed>.

5. WORK PLAN

5.1. T1. DATA COLLECTION AND PUBLICATION

5.1.1. Summary

This task will expand and consolidate the process of acquisition, long-term conservation, and publication of traces of EGEE activity and applications that has been successfully initiated by the Grid Observatory cluster in EGEE-III.

Type of activity: Service Activity (SA)

Effort: 5 FTE

Funding source: **Commission 3 FTE**; participating NGIs 3 FTE

Leading partner: LRI-CNRS

5.1.2. T1.1 Data acquisition

The primary role of the GO SSC is acquisition, curation, and long-term conservation of the monitoring data produced by the EGEE grid about its own behaviour. The SSC will continue its approach of building on the rich ecosystem of monitoring tools developed in gLite and by the users community, as well as the operations team with Nagios deployment. The GO SSC will thus limit its activity to exploiting their results, with one notable exception. Exploiting the results will take two paths

- Enabling the general deployment of the acquisition tools prototyped in the GO cluster of EGEE-III cluster, as a certified component of the gLite middleware.
- Long term conservation of the monitoring data collected by HEP experiments, currently gathered at at CERN, which are currently discarded after their immediate operational use is passed.

Both activities involve active collaborations with EGI-operations, and the first one involves collaboration with EMI.

This subtask represents a change of scale with respect to the GO cluster activity. The SixSq. Company will provide its experience with middleware packaging, and in interaction with the CERN operation team.

<2be developed, SixSq>

5.1.3. T1.2 Power consumption data

The notable exception quoted above is the acquisition of data related to power consumption. Due to limited access to such information, the research in optimization is often limited, with researchers often focused on a small-scale sub-problem that could be simulated.

With extended time-series of activity, the GO SSC will offer a ***unique and high visibility facility***.

It will allow to better understand the performance and energy trade off in the e-science area. This might be a point of particular interest for interaction with the cloud community, for which power consumption is a major issue: elucidate whether the power consumption of e-science application is comparable with the business workload of data centers.

5.1.4. T1.3 Data organization and presentation

<2be written –strongly dependent on partnership>

This task is an integration task for the whole project: acquisition as well as ontology is required to structure data description; results may be exploited in the information service organization task.

The task will first re-structure the datasets according to standards, either event-oriented or resource-oriented, for which standards exists or are in progress. This corresponds to "lossless" compression.

Two seconds steps must be considered in parallel:

- Data curation, which is presented in T3.2. Another source for this activity is the set of filters developed in gStat. This will also contribute to task 2.3, as a contribution to better accounting.
- Providing analysis facilities, ranging from basic statistics to the exploitation of stabilized analysis methods. The Matlab facilities on EGEE will be expoited.

5.1.5. T1.4 The GO portal

The GO portal is the visible part of the project. In EGEE-III, the GO portal has been built as a data repository.

Dynamic data presentation in response to data organization changes <2 be written HG>

Analysis methods will be proposed as on-line tools. <2 be developed – Logica ?>.

The GO portal is currently operated by the HG company. In the development of the project, the alternative of re-internalization, with HG as an expert, will be considered.

5.2. T2 CONTRIBUTIONS TO A SUSTAINABLE OPERATION MODEL

5.2.1. Summary

This task contributes to the overall grid management and governance goals towards sustainable, reliable and secure grid platforms. It will define advanced methods for analyzing, representing, and correlating the data. It will explore operational on-line usage of these methods, in order to evolve towards an *information provisioning service organization*.

Type of activity: Joint Research Activity (JRA)

Effort: <2 be defined>

Funding source: **Commission 1.5 FTE**[2]; participating NGIs; industry partners; collaborating projects.

Leading partner: UNIPM

5.2.2. T2.1 Models of the grid dynamics

<2be written>UNIPM + LRI

5.2.3. T2.2 Towards an information service organization

Information Technology (IT) Governance focuses on performance of IT systems and risk management. Industrial governance standards are captured in ISO 38500[3], which strongly focuses on managing the IT resources on behalf of stakeholders who expect a return from their investment.

Logica, an industrial partner that is experienced in the discipline of IT Governance, will lead this task. Its role will be to steer the interaction between the GO SSC and EGI-operations towards the application of IT governance standards and approaches.

The specific focus of this sub-task is **governance support by intelligent monitoring and learning agents.** Given the complexity of grid infrastructures, automated support to processes of information retrieval and analysis has become necessary. As explained supra, an extensive monitoring infrastructure does exist: gLite logs, and user-level software (eg HEP experiments), as well as the generic probing environment Nagios. We thus focus on the exploitation of the output of these monitoring tools, in an operation-oriented perspective.

[2] 2 FTEs would allow for easier relations with EMI and EGI-operations, by properly acknowledging their contribution

[3] ISO/IEC 38500 Corporate governance of information technology, (very closely based on AS8015-2005) provides a framework for effective governance of IT to assist those at the highest level of organizations to understand and fulfill their legal, regulatory, and ethical obligations in respect of their organizations' use of IT. ISO/IEC 38500 is applicable to organizations from all sizes, including public and private companies, government entities, and not-for-profit organizations. This standard provides guiding principles for directors of organizations on the effective, efficient, and acceptable use of Information Technology (IT) within their organizations.

Whereas the acquisition and interpretation of monitoring within individual domains is done superfluously, correlation between domains is not commonly done yet. In collaboration with subtask 2.1, we develop methods to correlate event-information between sites. We research how to automate the retrieval of application-level metrics. We will demonstrate tools that allow feeding back the results of these metrics through both automatic and administrative means to the site operations. Of primary interests are automatic feed-back loops that enable near-real time failure identification and remediation.

From the technical point of view, we intend to develop Nagios plugins that implement such functionality. A challenge thereby is that we have to take into account that Nagios plugins are not statefull.

From the organisational point of view, we want to focus on both, administration as well as user perspectives. Having generic EGI goals in mind, we therefore foster relationships with other SSCs, such as for example the LifeScience SSC.

The interaction with EGI-operations and EMI with ensure that the tools may be deployed on the live infrastructure.

5.3. T3 GRID ONTOLOGY

5.3.1. Summary

Concepts describing not only resources, as in the Glue Information Model, but also the Grid inputs (users, virtual organisations and applications), the Grid dynamics (lifecycles of individual jobs and the data transfers), applications to data consolidation and semantic inference.

Type of activity: Joint Research Activity (JRA)

Effort: 3 FTE

Funding source: **Commission 1 FTE**; participating NGIs 2 FTE

Leading partner: MIS-CNRS

5.3.2. Presentation

Task 3 deals with the building and exploitation of a grid ontology accounting for both structural and dynamical aspects of grids at different levels of abstraction. For the construction of the ontology, many resources are used as data: (i) existing termino-ontological resources on grids (GLUE, which is the basis for interoperability between the EGEE grid infrastructure and other grid infrastructures e.g. the Open Science Grid project in the US) will be considered as a main reference resource); and (ii) results from the modeling of grid dynamics.

5.3.3. Task 3.1 Ontology building

Goal

The goal of this task is to transform and enrich the GLUE, which is expressed as UML model, into an ontology based on logical descriptions of concepts in order to carry out inferences. The ontology will cover concepts already present in the GLUE (physical resources, components and services) but it will also include concepts about logical resources, jobs and their lifecycle and, generally speaking, the dynamics of EGEE, all kinds of concepts needed to reason on the traces.

Tools

Resources used are the foundational ontology DOLCE [Mas03], the core ontology of programs and software COPS [Lan09], and existing grid ontologies (covering mainly structural aspects).

Method

The informal descriptions associated to the entities and relationships structuring the conceptual model GLUE v. 2.0 are modeled in order to get a more formal and semantically richer model than the actual class model in UML. First, using DOLCE and the concepts coming from other termino-ontological

resources will enable restructuring the concepts coming from GLUE v. 2.0. Secondly, knowledge about the dynamics of EGEE are modelled into a semi-formal ontology and then a formal one, according to the OntoSpec methodology defined by MIS [Kas05].

5.3.4. Task 3.2 Inferences fo trace analysis and publication

Goal

The goal of this task is to develop and/or adapt a tool for semantic analysis of the traces. This tool will exploit the ontology to carry out inferences on grid traces (especially to detect inconsistencies), and also to improve information retrieval tools for task 1 and task 3.

Tools

A set of tools is needed to manage and efficiently access the ontology, and to carry out inferences on traces. Inferences will be carried out on semantic representations of traces, so a tool is required to built such semantic representations from log files. Several semantic engines exist which are currently used in numerous projects. The choice of tools will be done in the beginning of the second year of the project.

Method

The choice of tools will depend on assessment of the number of data and concepts to be taken into account and also of the database technologies that will be chosen. Large tests of chosen tools in trace analysis will allow to validate the ontology and to improve it. The tools will be extended to link them to the publication tool of the task 1.

5.4. T4 FOSTERING SCIENTIFIC COLLABORATION

5.4.1. Summary

Proactively engage the relevant scientific communities and institutions, create and implement the interaction processes, contribute to on-line analysis tools.
Type of activity: Networking (NA)
Effort: <2 be defined>
Funding source: **Commission 2.5 FTE,** Participating NGIs; collaborating projects
Leading Partner <2 be defined>

5.4.2. T4.1 Engaging the Autonomic Community

This action will target primarily the NSF center for Autonomic Computing, and the corresponding European networks.

The focus of this task is the cross-fertilization between AC and production grids. The first axis is the need to demonstrate the effectiveness of autonomic concepts for end users who utilize Grid computing infrastructure. The major issue is to define a process that could be an equivalent of undertaking the "*paper-in-Nature*" test, whereby a scientist could claim that through the use of autonomic computing techniques they were able to reach a particular scientific insight that they could not have obtained otherwise. The process must focus on effective dissemination of such benefit to get community commitment. There are actually multiple dimensions to this issue: (i) demonstrating a reduction in cost when using the concepts; (ii) demonstrating that there was added value above existing tools; (iii) demonstrating improvements in QoS, with a particular focus on how such QoS leads to improvements in the scientific findings.

It must be noted that the interaction with AC concepts covers the whole range of grid exploitation, from infrastructures to applications and including data quality. One of the key challenges to improve take up of autonomics was the need to empower end users and developers to address uncertainty. On the other hand, due to the very nature of problem solving approaches in autonomic computing, it is necessary to attempt solutions of specific problems, before attempting to generalize.

To be practical in addressing this issue, we propose to constitute a **knowledge base of *Open Issues in Grid Operations and Applications*.** Many of them fall in the scope of Autonomic Computing. The problem will be specified informally, but with formal performance criteria to be met. For operation entries, the corresponding entry in the knowledge base will include relevant and significant datasets and when possible, documentation. For application entries, the publicly accessible Grid applications will be exploited to assess issues such as scalability and performance for the use of autonomic computing algorithms. Such problems, operation and application-level, will become a basis for scientific challenges, or AC benchmarks. To make sense, the process has to be two-ways, and iterative, meaning that it has to include a process for specifying, and guarantying the significance of the datasets, as well as evaluating the potential solutions or hints provided by the AC community.

5.4.3. T4.2 Engaging the Distributed systems community

<2be written>

- Relations with Core Grid
- Evaluating the feasibility of creating an experimental testbed for gLite. This is in relation with possible collaboration with the G5K infrastructure in France.

5.4.4. T4.3 Interactions with OGF

<2be written>

- Ontology
- Scheduling

5.5. T5. MANAGEMENT:

5.5.1. Summary

Coordinate the efforts between national, regional and European level. Manage the SSC and be involved in EGI.eu boards.
Type of activity: Management (MGT)
Effort: 0.5 FTE
Funding source: Leading participating NGI

This subtask has in charge the coordination of the whole GO cluster.

- Manage interactions with the EGEE/NA4 and with the EGEE Project Office, including administrative aspects, which are not familiar to some participants.
- Monitor and coordinate activities of all other subtasks and the progress for each of them.
- Set up tools and services to foster the collaborative work within the SSC (mailing lists, wiki, documentation repositories, etc.). A specific issue is the relationship between publicly accessible collaborative tools associated with the web portal, and the tools internal to the project.
- Organize cluster meetings, both face-to-face and remote (through collaboration tools).
- Produce reports and other documentation requested by EGEE/NA4 SC and by the Project Office.

5.6. BUDGET AND STRUCTURE

Required effort from the from the European Commission: 8 FTEs
TOTAL EFFORT: <2 be defined>

Organization

		T1	T2	T3	T4	T5
Funded partners	LRI	X, LP		x	X, LP	x
	MIS		X, LP			
	LPC		x			
	UNIPM			X, LP		
	IC	x				
	Logica			x		
	Health Grid	x				
	SixSq.	x			x	
	IBM	x				
	X CoreGrid		x	x	x	
	Y CoreGrid			x	x	
Networking partners	CAC				x	
	U. Bologna		x		x	
	U. Cardiff				x	
	EGI-operations			x		
	EMI			x		

6. NON-HUMAN RESOURCES

Given the relatively limited needs of the SSC, the most efficient solution is to require them from only one NGI. The French NGI will provide:

- Storage capacity. This includes disk storage (GRIF site) and backup facilities (CC IN2P3), as well as the generic software environment (database facilities).
- Computing resources, targeting exploitation of analysis tools. The GRIF site has experience with the gLite/Matlab interface, and will be the main supporting center for this activity.

The modalities of access to the CERN immaterial resources (datasets from experiences) are one of the activities of the SSC.

7. SUSTAINABLITY PLANS

The GO SSC is intended to be a stable entity whose primary goal is safeguarding and publishing datasets in the long run. Consequently, it is extremely important that the GO SSC can evolve towards a permanent structure and define a sustainable financing model. On the other hand,

- the activity is much younger than all other scientific SSC in the EGEE framework, and is still in it ramp-up phase; thus it requires initial development funding;

- generally speaking, the computer science community heavily relies on national and international project-oriented funding schemes.

The plan is thus as follows.

- The French NGI ensures the infrastructure requirements, including non-dedicated operation costs, together with EGI, (e.g providing the volatile and long-term storage). The national production grid has been recognized as a TGIR (Très Grande Infrastructure de Recherche) by the French ministry of research.
- A mix of permanent and project-based funding will provide human resources.
  - Permanent scientific personnel will provide scientific and operational steering.
  - The major source of project-based funding for this period will be, on one hand the EC (through this call, and possibly a COST programme), on the other hand the national or regional funding schemes (eg the French ANR, the UK e-science programme).
- Industry partners are now involved. This will facilitate applying for national or international funding scheme targeting industry R&D; in a more ambitious scheme (such as the evolution of the TOP500), the GO portal would become sufficiently popular to attract company sponsoring.

It is of major importance to define the perimeter of consolidation for an SSC. In France (and in some other countries as well), the combination of the administrative system, and the FP7 administrative rules, preclude any possibility to consolidate 1) FP7 funding, and 2) national project-based funding, at the accounting level. In practice, matching effort can come only from permanent staff, or temporary staff hired on recurrent funding. Thus, the EC should precise the guidelines for evaluating the evolution towards sustainability, which cannot be based solely on accounting reports.