EGI-InSPIRE

LIP-IBERGRID NGI OLA QUESTIONNAIRE

Document identifier:	CSIC-IBERGRID-EGI-InSPIRE-OLAs-NGI questionnaire-v1.doc
Date:	17/06/2010
Activity:	**SA1**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/58

NGI OLA QUESTIONNAIRE

Date: 17/06/2010

PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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

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

TABLE OF CONTENTS

1.OLA NGI QUESTIONNAIRE

1.1.OLA STATUS

1.Number of certified sites in the NGI

Spain has 18 certified sites but it is working inside the NGI_IBERGRID with a total of 26 certified sites (information from GOCDB 17/08/2010).

2.Number of sites that have already signed an OLA or comparable document

11 sites certified during EGEE already signed the EGEE SLA (http://www.egee.cesga.es/EGEE-SA1-SWE/accounting/sld.html). The new sites certified sites certified during EGI has not signed any document yet. .

3.In case of a comparable document being used, describe deviations from the metrics used in the original EGI OLA document.

EGEE SLA was the only document used until now. The previous document has basically the same requirements for sites that the new EGI OLA.

4.What is the main obstacle to the adoption of the OLA by all sites?

The current OLA (https://documents.egi.eu/public/ShowDocument?docid=31) seems to be identical to the EGEE SLA, and in this sense, there should be no major obstacles in adopting it in IBERGRID. The only minor issue is that some of the small sites within IBERGRID, when they signed the EGEE SLAs, have requested to decrease the amount of minimum storage space they were supposed to offer, and be more flexible in the time frames for handling GGUS tickets.

5.Which are the main considerations / objections of sites to the OLA?

Some of the small sites within IBERGRID, when they signed the EGEE SLAs, have requested to decrease the amount of minimum storage space they were supposed to offer, and be more flexible in the time frames for handling GGUS tickets.

6.Describe any modifications that you would consider to the OLA metrics definitions?

Four working hours to acknowledge GGUS tickets could be too short. We would propose to extend that metric to next working day.

7.Are there any metrics that should be added/removed from the OLA? Include a brief justification for your answer.

We would add to NGI obligations a metric (or a recommendation) for (at least) one operational meeting per month with sites. These meetings are normally good practice where sites can share experiences, communicate problems and receive directives from NGI managers. On the other side, We would also add a recommendation for sites to participate in those NGI meetings, at least, once per month.

1.2.ENFORCEMENT METHODOLOGY

8.Are there any improvements you would propose to apply in your NGI to the current enforcement methodology of the OLA? (Monthly League Table, justifications for breach of A/R metrics)

If there is a metric for NGI operational meeting (as we proposed in question 7) one could also enforce the participation of underachieving sites in the first NGI operations meeting after the A/R publishing. In this way, NGI managers can also check the issues which lead to such bad performance, and try to educate the site in such a way that the situation does not happen again.

9.What kind of rewards/penalties for sites would you consider for over/underachieving sites?

The policy that underachieving sites (for 3 consecutive months) may be suspended seems OK. Best sites should also be distinguished (for example, a TOP-10 table of the best sites distributed within the Monthly League Table).

10.Do you find the current system for providing justifications for A/R failures adequate? If not why? What else would you use?

We think opening tickets to sites via GGUS is adequate.

11.Do the justifications in general adequately describe the incident, main cause and the recovery strategy used?

In general, the justifications (for incidents) provided by site administrators are not self-explanatory.

1.3.MONITORING TOOLS

14.Describe any defects that you've encountered with the OLA monitoring tools currently used (e.g. Nagios, GridView)?

There are some issues with GridView about collecting the A/R metrics. The issue was opened by IBERGRID staff and is tracked via https://gus.fzk.de/ws/ticket_info.php?ticket=60594

The transition to NAGIOS is not straight forward to sites since it is difficult for sites to get details about NAGIOS failures. Those details should be available through MyEGEE portal, but that portal still presents several issues, like the ones described in https://gus.fzk.de/ws/ticket_info.php?ticket=60597 and https://gus.fzk.de/ws/ticket_info.php?ticket=59403. Moreover, site admins complain that it is not easy to understand what NAGIOS tests are doing, and have access to script samples to reproduce the problem in a control environment.

15.Describe any improvements that you would consider to the OLA monitoring tools currently used (e.g. NAgios, GridView)?

I think the A/R metrics do not take into account the size of the site, namely the amount of computing power and storage power it is offering. It is different to compare A/R values for a 8 CPUs / 1TB site with the results of a 8000 CPUS / 1000TB site. Therefore, some proper normalisation to the A/R results has to be found in the near future.

In a more general way, the navigation via the current GridView page is complicated. Just trying to get availability information for one of my sites, I've seen that, for example, I have to search my site name

from a huge list of sites. I should be able just to search for my site name. Also, most of the navigation is aimed for wLCG Tier-sites which is, from the infrastructure point of view, a wrong approach.

1.4.FUTURE DEVELOPMENTS

16.Do you think that the OLA should remain part of site certificate process or there is a different procedure you would like to use?

No. When sites join the infrastructure, the staff is not fully aware of the proper procedures and have to be "educated". Therefore it is nonsense to force sites to sign an OLA that they may not be able to fulfil. This does not mean that the EGI OLA criteria shouldn't be applied to them. They should try to get good A/R values and explain if they don't manage to do so. However, a six month grace period could simply consist on not being suspended if they fail to be above the A/R threshold for 3 consecutive months. After the grace period, they have already some experience, and know what they are dealing with, and therefore, are in a much better position to accept and sign the proposed OLA.

17.How do you (or would you) manage OLAs in your NGI?

OLAs were still not discussed internally but most probably we will adopt the EGI proposed OLA with changes where we think appropriate.

18.Would you object to an increase of the minimum Availability/Reliability thresholds to 80% and 85% and respectively?

At this transition point yes. However, as the infrastructure stabilizes, I could consider that change in the end of the 2nd year of the project.

19.Would you object to permitting a grace period of 6 month for new sites were availability and reliability thresholds are 70% and 75% respectively?

As referenced in question 19, I would consider the same A/R along the first two years of the project, and then, increase those thresholds if appropriate. The grace period for sites would consist on not being suspended if they fail to be above the A/R threshold for 3 consecutive months.

20.What thresholds would you like to see for EGI core servicers? Do you agree with 80%/85% as in sites?

Since core services are supposed to be deployed by more experienced sites, and are critical from the infrastructure point of view, I agree with a higher A/R value for them.

21.Please provide any additional comments that were not covered with the previous questions

I wonder how the enforcement of OLA will be pursued and if benchmarking will play an important role to determine the computing and storage power that sites are offering to the infrastructure. In this particular aspect, we are a bit concern with the adoption of HEPSCPEC06 which has been "elected" as the standard benchmark to normalize CPU usage in EGEE. HEPSCPEC06 does not represent the computing environment for a multiscience platforms, and it is only representative for HEP applications. Is it going to be one of the requirements in the OLAs?

The OLA document is written in terms of NGIs. However, in practice we have several NGIs operating in federations. Therefore, section 3.1 should be re-written alerting to that fact that, for operations, a "NGI body" could represent/contain several NGIs. For example, NGI_IBERGRID is the operational body that represents Spanish and Portuguese NGI.