

EGI-InSPIRE

Polish NGI - PL-GRID OLA QUESTIONNAIRE

Document identifier:	EGI-InSPIRE-OLAs-PLGrid-v0_4
Date:	17/06/2010
Activity:	SA1
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/58

Copyright notice:

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration.

EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years.

This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”.

Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders.

The information contained in this document represents the views of the copyright holders as of the date such views are published.



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.....	4
1.1.OLA STATUS.....	4
1.2.ENFORCEMENT METHODOLOGY.....	4
1.3.MONITORING TOOLS.....	4
1.4.FUTURE DEVELOPMENTS.....	4

1. OLA NGI QUESTIONNAIRE

1.1. OLA STATUS

- Number of certified sites in the NGI
9, providing ~10k CPU cores
- Number of sites that have already signed an OLA or comparable document
7 sites signed EGEE SLA which we consider still binding until we sign new OLA.
- In case of a comparable document being used, describe deviations from the metrics used in the original EGI OLA document
It is not actually used yet, but PL-Grid is preparing an OLA which is based on EGEE SLA with extensions towards resource allocation i.e.:
 - volume of resources available for resource allocation process supported by the NGI**
 - parameters for responsiveness during resource allocation process.**
- What is the main obstacle to the adoption of the OLA by all sites?
It is not clear yet, we expect to have more experience on that matter when our will start to enforce new OLA.
- Which are the main considerations / objections of sites to the OLA?
It seems it is too early for that question, see also point 4 above.
- Describe any modifications that you would consider to the OLA metrics definitions?
The metrics proposed and defined in EGI OLA will be adopted by Polish NGI for sake of keeping consistency between EGI and Polish NGI. However, in EGI OLA we would like to propose changes (or clarifications) towards better definition of site responsiveness in terms of GGUS tickets, security incidents handling, requests for middleware upgrades.
Issues with the current proposed metrics:
 - Maximum time to resolve GGUS incidents**
If we take ITIL definition of incident it reads: “unplanned interruption to an IT service or reduction in the quality of IT service”. So any problem causing an interruption falls under this definition. Our experience shows that time to resolve it depends on a problem type. If there is a recipe which can be applied in a particular case then we can give some estimation

on the solution time otherwise it may be impossible to find a solution within proposed 5 business days.

We suggest that if there is no known or applicable recipe for a particular incident a suitable procedure to handle the ticket should be applied and no OLA breaking should be accounted such a case. Such tickets should be under special attention as they may indicate fields that needs improvements from the project side.

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

a) Core services (WMS, VOMS, LFC etc.) availability/reliability metric are needed. Core services are even more important than resources at sites (without working VOMS nobody can use site's resources).

b) New metric is needed depending on how do we define GGUS incidents. If security incidents will be handled by GGUS tickets then we may reuse GGUS incident handling rule, but is that consistent with security rules of not exposing such cases to wide public (as GGUS tickets are)?

A separate metric describing site reaction to security incidents will be useful.

c) Middleware upgrade usually make a problem in terms of keeping the sites at some recent release. Thus a metric describing site reaction in terms of middleware upgrade would be useful. It could say that site has 3 months for the release.

d) We suggest introduction of “yearly availability” metric and lowering minimum monthly metric. Both metrics could be accounted monthly. The point is to reflect the real operation of a typical computing center. Sometimes it requires to put the site under 2-week maintenance. This would lead to 50% monthly availability decrease. However, such works do not happen often. If a site declares 80% of yearly availability and 50% of monthly availability it allows to do 3 big maintenance works during the year and keep the remaining months at 90% at minimum.

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)

For a long term the OLA reports should be easily analysable in some on-line tool where details and requests for corrections can be entered. We believe that the ability to analyse and report discrepancy between figures observed at site and reported in OLA would greatly improve acceptance of sites for OLAs. The monitoring staff should be able to approve/reject reports of monitoring system failures affecting the OLA figures.

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

Our understanding of OLA is that it defines entry level to the infrastructure. Rewards are not necessary, however history of the metrics could be interpreted as confirmation of solid operation (useful for e.g. VOs). Some smart ranking of sites based on long term data would be useful as well. The only way we could think of rewarding an NGI would be to decrease the EGI membership fee ;-) It would even have a practical justification as the well operating NGI should generate less effort for the EGI than others. However, our observation is that small, rarely used sites are able to achieve very high metrics so a weight for site usage should be applied here.

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

The system with GGUS ticket is fine, but a procedure for taking the justification into account need to be defined.



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

We feel that the OLA, OLA monitoring system and metrics should be designed in a way that does not require any justification but penalty action is executed when thresholds not met (e.g. site suspension). However, before that, the process should take into account “request for figures amendment” when e.g. site admin disagrees with reported values. We see a space needed for reporting some general problems with not meeting many metrics asking for not executing penalties if the reasons were not due to bad operations but other. Such a “reprieve request” should be send to OCC or similar for consideration.

1.3. MONITORING TOOLS

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

It would be good to clarify what happens when NGI monitoring system fails. We see solution of “last recent known status taken as valid” as insufficient. Such approach is unfair for sites which failed at the moment as they most probably will be trying to change the state. If they succeed and the monitoring is still down their figures will be incorrect.

Availability computation should be extended to embrace cases of sites state change e.g. certification, suspension.

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

a) Service levels should be customizable per site and service. Common thresholds does not reflect the reality, instead individual site should have there own thresholds per service (site, WMS, etc.) .

b) In EGEE-nagios monitoring system the NGI is responsible for hosting the system and providing monitoring data. This allows for easy data distortion. Do we plan to check the integrity of published data? In general we feel that the NGI is not neutral here and should be a subject for validity checks.

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?

a) Our understanding is that signing OLA should be a part of site certification process. It is good to define the rights and responsibilities from the very beginning. However, we are not sure that formal signing of paper documents is necessary in each case. However in each case the NGI should be able to define thresholds which are valid for the NGI/sites. In general there should be no sites for which the OLA values are not known.

17. How do you (or would you) manage OLAs in your NGI?
The OLA is signed when the site joins the infrastructure. The site is then informed about rights and obligations. Then the OLA procedures are just executed.
18. Would you object to an increase of the minimum Availability/Reliability thresholds to 80% and 85% and respectively?
No, however please see “d)” comment for point 7 above. We would like to point out difference between Availability and Reliability. Availability can be affected by some factors the site is not able to mitigate. However, reliability requires just good site operation.
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?
Yes. Our experience confirms that 1 month is enough for a new site to solve all initial problems. OLA metrics thresholds for that site should be kept at target level, however the project can refrain from penalties for the site for the grace period of 1 month. Keeping the thresholds at the target level makes OLA management easier and allows to see what are the real, best values achievable by the site. If a site had problems to achieve them they should give feedback and based on this the EGI/NGI should improve procedures for introducing new sites. It would be nice to mark those sites in monthly reports that they are newcomers.
20. What thresholds would you like to see for EGI core servicers? Do you agree with 80%/85% as in sites?
Yes.
21. Please provide any additional comments that were not covered with the previous questions
a) Each site should be able to define their own thresholds and own thresholds per service. OLA monitoring tools should be adapted accordingly.
22. **b) OLAs between EGI and NGIs is important to define, it seems the two documents may be related. There are some responsibilities that NGI should provide for sites e.g.**
1. 1st line support team, 1st line response time
 2. tools etc.
23. **c) In the current model, it is possible to have a site in the infrastructure, that operates at 100% , but not provide any resources to users (e.g. has queues disabled). We think the latter is the crucial 'service' in the infrastructure, therefore resource allocation should be covered by OLA.**