

EGI-InSPIRE

OPERATIONAL LEVEL AGREEMENTS WITHIN THE EGI PRODUCTION INFRASTRUCTURE

EU MILESTONE: MS404

| | |
|----------------------|---|
| Document identifier: | EGI-MS404-v0.12.doc |
| Date: | 12/08/2010 |
| Activity: | SA1 |
| Lead Partner: | AUTH |
| Document Status: | FINAL |
| Dissemination Level: | PUBLIC |
| Document Link: | https://documents.egi.eu/document/65 |

Abstract

This document describes the status of Operational Level Agreements – OLA(s) – within the European Grid Infrastructure. It illustrates the current adoption status of existing OLA(s) within the National Grid Infrastructures, and defines plans for OLA(s) extensions to enhance the quality of the pan-European Grid infrastructure and its services within the EGI-InSPIRE project.

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.

Delivery Slip

| | Name | Partner/Activity | Date |
|--------------------|--|-------------------------|-------------|
| From | | | |
| Reviewed by | Moderator: Reviewers: | | |
| Approved by | AMB & PMB | | |

Document Log

| Issue | Date | Comment | Author/Partner |
|--------------|-------------|--|--|
| 1.0 | 18/6/2010 | First TOC | Dimitris Zilaskos, Christos Kanellopoulos |
| 2.0 | 25/6/2010 | Revised TOC | Dimitris Zilaskos, Christos Kanellopoulos, Tiziana Ferrari |
| 3.0 | 1/7/2010 | First draft | Dimitris Zilaskos, Christos Kanellopoulos, Tiziana Ferrari |
| 4.0 | 7/7/2010 | Second draft v 0.8. Internal review | Tiziana Ferrari |
| 5.0 | 29/7/2010 | Final version 0.12 addressing internal review comments | Dimitris Zilaskos, Christos Kanellopoulos, Tiziana Ferrari |
| 6.0 | 12/8/2010 | Version 0.13 addressing additional comments from S. Newhouse | Tiziana Ferrari |

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. INTRODUCTION | 5 |
| 1.1. PURPOSE | 5 |
| 1.2. APPLICATION AREA | 5 |
| 1.3. REFERENCES | 5 |
| 1.4. DOCUMENT AMENDMENT PROCEDURE | 6 |
| 1.5. TERMINOLOGY | 6 |
| 2. EXECUTIVE SUMMARY | 7 |
| 3. DEFINITIONS | 8 |
| 4. CURRENT OLA(S) | 9 |
| 4.1. SITE-NGI OLA | 9 |
| 4.1.1. <i>Metrics and related measurement tools</i> | 10 |
| 4.1.2. <i>Enforcement procedure</i> | 11 |
| 4.2. 1 ST LINE SUPPORT OLA | 11 |
| 4.2.1. <i>Metrics and related measurement tools</i> | 12 |
| 4.2.2. <i>Enforcement procedure</i> | 12 |
| 5. ROADMAP | 12 |
| 5.1. EXTENSIONS OF EXISTING OLAS | 12 |
| 5.1.1. <i>Tuning of thresholds</i> | 12 |
| 5.1.2. <i>Cases for site suspension</i> | 13 |
| 5.1.3. <i>Differentiated Quality of Service</i> | 13 |
| 5.1.4. <i>New services</i> | 14 |
| 5.2. NEW OLA(S) | 14 |
| 5.2.1. <i>NGI-EGI OLA</i> | 14 |
| 5.2.2. <i>Site-VO OLA</i> | 16 |
| 5.2.3. <i>EGI-VRC OLA</i> | 16 |
| 5.2.4. <i>Tool extensions</i> | 16 |
| 6. CONCLUSIONS | 16 |
| 7. ANNEX A: NGI QUESTIONNAIRE | 18 |
| 7.1. OLA STATUS | 18 |
| 7.2. ENFORCEMENT METHODOLOGY | 18 |
| 7.3. MONITORING TOOLS | 18 |
| 7.4. FUTURE DEVELOPMENTS | 18 |

1. Introduction

1.1. Purpose

This Milestone presents the status of Operational Level Agreements currently used in the production infrastructure developed during the EGEE-III project, and elaborates on future extensions that will be developed during EGI-InSPIRE that are needed to better reflect the current 1st line support structure within EGI, the quality of service to be provided to Grid end-users by EGI resource centres, the response time of EGI site managers to trouble tickets, and generally speaking the quality of service of EGI and NGI core software and operational services.

1.2. 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.

1.3. References

Table 1: Table of references

| | |
|-------------|---|
| R 1 | Information Technology Infrastructure Library web site http://www.itil-officialsite.com/home/home.asp |
| R 2 | EGEE III SA1 SLA https://edms.cern.ch/document/860386 |
| R 3 | Operational Level Agreement between GGUS and TPMs within the ROCs https://edms.cern.ch/document/888089 |
| R 4 | Site-NGI OLA https://documents.egi.eu/document/31 |
| R 5 | GridView https://twiki.cern.ch/twiki/bin/view/LCG/GridView |
| R 6 | Nagios for Grid monitoring https://twiki.cern.ch/twiki/bin/view/LCG/SAMToNagios |
| R 7 | GGUS tool https://gus.fzk.de/pages/home.php |
| R 8 | GOCDB tool https://goc.gridops.org/ |
| R 9 | Middleware in EGI http://knowledge.eu-egi.eu/knowledge/index.php/Middleware Components and Middleware Consortia |
| R 10 | EGI League Tables https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics |
| R 11 | GGUS metrics reports https://gus.fzk.de/pages/metrics/download_metrics_reports.php |
| R 12 | GGUS report generator https://gus.fzk.de/stat/stat.php |
| R 13 | Nagios http://www.nagios.org/ |

| | |
|-------------|---|
| R 14 | EGI Availability/Reliability results https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics |
| R 15 | EGEE III User Support Advisory Group https://twiki.cern.ch/twiki/bin/view/EGEE/SA1_USAG |
| R 16 | Downcollector https://ccenoc.in2p3.fr/DownCollector/ |
| R 17 | Service Level Agreement http://en.wikipedia.org/wiki/Service_level_agreement |

1.4. 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>

1.5. Terminology

A complete project glossary is provided in the EGI-InSPIRE glossary:
<http://www.egi.eu/results/glossary/>.

2. EXECUTIVE SUMMARY

Operational Level Agreements (OLAs) define how IT groups work together to meet IT service level requirements and are defined in the Information Technology Infrastructure Library (ITIL) [R 1], whose purpose is to optimize the delivery of IT services to customers and users. The OLA is an internal “back to back” agreement that defines how two different organizations will work together to support the delivery of defined IT services to customers and users. Within EGI, OLA(s) serve different purposes:

- to ensure mutual understanding of the principles of cooperation between EGI parties;
- to define the responsibilities of each party;
- to set the procedures for monitoring the fulfilment of commitments towards the users of the infrastructure and of the operational services;
- to define a set of requirements that satisfy the users and the operators of the infrastructure;
- to establish reporting and problem-solving procedures.

OLAs inherited from EGEE will continue to be enforced and will evolve to meet the new requirements that are expected within a distributed de-centralised infrastructure. EGI OLA(s) will be the pillars needed to develop a fully service-oriented sustainable infrastructure.

Additional OLA(s) needed to cover the new interactions between the EGI partners will be identified and produced with the aim of maintaining a uniform smooth experience for users of the infrastructure across the growing numbers of participating NGIs. These OLA(s) could include:

- EGI to NGI Core services OLA (e.g. EGI Certification Authority)
- EGI to NGI Central Operational Services OLA (e.g. Central GOCDB)
- NGI to EGI Core services (e.g. WMS, Accounting infrastructure)
- Differentiated Quality of Service (e.g. for sites and VOs with different Availability/Reliability requirements)

3. Definitions

| | |
|-----|---|
| SLA | A "Service Level Agreement" usually refers to a formally negotiated agreement between two parties. It is a contract that exists between customers and their service provider, client or between service providers. It records the common understanding about services, priorities, responsibilities, guarantee, and collectively, the level of service. For special restrictions on the use of this term in EGEE, please see the "background" section of this document. |
| SLD | Service Level Description (SLD) is the name given to the initial EGEE-II ROC-Site agreement. |
| OLA | An Operational Level Agreement (OLA) defines the interdependent relationships among the internal support groups working to support a Service Level Agreement. The agreement describes the responsibilities of each internal support group toward other support groups, including the process and timeframe for delivery of their services. The objective of the OLA is to present a clear, concise and measurable description of the service provider's internal support relationships. |
| NGI | A National Grid Initiative or a European International Research Organizations that is part of the EGI infrastructure |
| ROC | A "Regional Operation Centre" is one of the Regional Federations in which EGEE was divided. |

4. Current OLA(s)

The EGEE project series produced two OLAs, one describing the services and relationship between resource centres and ROCs [R 2], and one defining the duties and quality parameters of 1st line support [R 3]. These OLAs define a set of metrics that cover some of the most crucial aspects of a pan-European production Grid infrastructure, that being site performance and reaction to support requests issued through the helpdesk. These OLA(s) serve as a basis in the EGI start-up phase to ensure the smooth continuation of the operations from EGEE to EGI.

OLA parameters – in particular site availability and reliability and ticket assignment and response time – are subject to periodic measurements and procedures were put in place to ensure that those quality parameters are met. In what follows we provide an overview of the existing EGI OLAs.

4.1. Site-NGI OLA

Improving the quality of the infrastructure has been an ongoing task throughout the EGEE project series. To that end, effort was placed in raising the awareness of Regional Operations Centres (ROCs) and resource providers about the importance of introducing measurable Key Performance Indicators (KPIs) and minimum quality targets were defined. These targets were made part of the SLA between a site and the corresponding ROC and initially it was planned that the SLA would be signed by every site participating in the EGEE production infrastructure.

The adoption of SLAs has been progressively expanding with the signing of the SLA and the testing of site conformance becoming part of the site certification procedure. Nevertheless, the global adoption of the SLA is usually difficult because of the legal implications in certain legal jurisdictions, and because of the lack of penalties.

In order to facilitate acceptance of the SLA where the term carried legal implications, the term SLD was used as synonymous, along with the addition of an explicit mention in the SLA text that it is not legally binding. Despite these measures, concerns were not fully addressed, and also additional concerns were raised to the point of the whole exercise.

We propose to replace the EGEE SLD terminology with “Site-NGI OLA” according to the ITIL standard. An EGI updated version of the Site-NGI OLA [R 4] is now available in the EGI document repository.

The Site-NGI OLA provides a skeleton and defines a fixed set of quality parameters and the minimum performance threshold expected. Whilst the minimum set of parameters and the corresponding thresholds are not supposed to be modified, still the Site-NGI OLA can be customized to adapt it to the local needs. The OLA defines the availability, reliability and support metrics on which the sites will be evaluated, the responsibilities of both parties, as well as the targets that must be met for each metric.

The site-NGI OLA currently defines:

- that baseline middleware versions are adopted by the resource centres [R 9];
- the minimum amount of resources and grid services to be provided, as defined in table 2 (see page 10);
- the metrics measured on a monthly basis through automated report generator tools – mainly GridView [R 5] to gather test results from the Nagios-based EGI monitoring infrastructure [R 6]). Statistics take into account site feedback.

The NGI responsibilities towards a site are also covered. The NGI has to provide:

- a help-desk facility for operational support sites;
- 1st and 2nd line operational support to site managers. Tickets are monitored to make sure they have been appropriately addressed by the site administrators and support staff
- the registration of site administrators into its support facilities.

Similarly, resource centres should:

- adhere to the operational procedures, which define security, operational and availability policies;
- ensure that relevant information in GOCDB is accurate [R 8];
- run supported middleware versions;
- respond to GGUS tickets swiftly;
- adhere to the terms and conditions of the OLA.

4.1.1. Metrics and related measurement tools

The metrics currently measured are summarized in table 2. Currently availability and reliability are calculated by gathering periodic monitoring results from the Nagios-based infrastructure. Availability and reliability are calculated by GridView on a hourly, daily, weekly and monthly basis.

The response times to tickets opened centrally through GGUS is monitored and reports are periodically generated [R 11] [R 12].

Table 2: Site-NGI OLA metrics

| | Value | Section |
|--|--------------------------|---------|
| Minimum number of site BDII's | one | 8 |
| Minimum number of CEs or SEs | one | 8 |
| Minimum number of WN CPUs/cores | eight | 8 |
| Minimum capacity of SE(s) | one TB | 8 |
| Minimum site availability | 70% | 10 |
| Minimum site reliability | 75% | 10 |
| Period of availability/reliability/outage calculations | per month | 10 |
| Minimum number of system administrators | one | 11 |
| Maximum time to acknowledge GGUS tickets | four hours | 11 |
| Maximum time to resolve GGUS incidents | five working days | 11 |
| Minimum number of supported user-community VOs | one | 11 |
| Tracking of SLA conformance | monthly | 12 |

The aforementioned metrics and quality thresholds define a functional grid site with an acceptable response time on handling operational issues.

Functional testing of the grid services exposed by a site is implemented through a Nagios-based monitoring infrastructure [R 13], which is also used to produce alarms in case of critical failures. Probes are executed against a site to simulate common workflows and tasks that a grid user performs. Successful execution of the probes indicates that a site is functioning properly. Nagios results are gathered from the NGI monitoring infrastructure through an ActiveMQ communication bus and are stored in a central database that is the data source for the calculation of the monthly statistics. In EGEE-III the Nagios probes were developed in the framework of the Operations Automation Team [R 6], whilst in the EGI era probes will be released by the third-party middleware providers whose components will be part of the UMD distribution.

4.1.2. Enforcement procedure

The EGI availability and reliability statistics are used to enforce the Site-NGI OLA. Publicly available reports [R 14] are an incentive for sites to achieve better results. Grid sites are requested to provide a minimum availability and reliability, and sites that fail to do so are requested to provide justification for their poor performance. In addition, suspension is imposed to site that fail to provide less than 50% availability for three consecutive months.

A new procedure for the gathering of justifications and to follow up sites in need of suspension is currently under experimentation. The EGI operations oversight team (also known as central COD) is responsible for chasing sites according to the following procedural steps.

- The results are produced by the GridView team in Excel and PDF format during first week of the month.
- Results are checked by TSA1.8 staff for any obvious errors and/or inconsistencies.
- The results are e-mailed to the NGI Operations managers mailing list and a ticket is opened to the COD Support Unit to trigger the site follow-up procedure.
- COD analyzes the results for underperforming sites and opens child tickets to the respective NGIs.
- Cases of underperforming sites that do not reply to tickets in seven days are escalated.
- Sites eligible for suspension are suspended by default upon proper notice to the site and respective NGI. The default suspension applies to all sites, unless the NGI requests the suspension to be held up. Holding up the suspension of a site is approved by EGI Chief Operations Officer under only exceptional and well motivated circumstances.

4.2. 1st line Support OLA

Purpose of the 1st line support OLA is to define the duties and response time to tickets of the 1st line support group within GGUS called Ticket Process Management (TPM). The 1st line support OLA was introduced during the EGEE-III project and clearly describes the parties involved in the GGUS system with the respective responsibilities and service hours. Goal of TPM is to quickly identify the nature of the problem described in the ticket and solve it if possible, or to assign it to the proper specialized support unit, and ensure that tickets are properly followed up until they reach a terminal state. TPM needs to respond to tickets or assign them to a different Support Unit within one business hour.

4.2.1. Metrics and related measurement tools

The GGUS tool has the ability to produce detailed reports about the tickets processing timelines [R 11]. These reports provide detailed information for many Support Units (including TPM), such as the number of tickets handled, the average solution time, response time and the number of tickets unsolved. In addition, TPM ticket escalation reports are generated on a weekly basis reporting about the number of tickets assigned, the ticket routing process (number of hops) and the reaction time [R 12]. Important metrics are:

- Ticket Age: number of days since the ticket was opened.
- Inactivity index: a score produced by an algorithm taking into account the ticket age, and the time passed since the last update was made by TPM or a support unit. High scores are an indication of a ticket that was not handled properly.

4.2.2. Enforcement procedure

In EGEE-III enforcement of the 1st line support OLA was performed during the User Support Advisory Group (USAG) meetings where the escalation reports were evaluated [R 15]. The reports were also submitted to SA1 management for evaluation.

In EGI a new procedure needs to be defined to streamline the followup process. A procedure similar to the one adopted for the Site-NGI OLA could be adapted, where the escalations reports are emailed to support unit leaders and the SA1 management. For each Support Unit thresholds of inactivity index and response time need to be defined depending on the criticality of the tickets.

5. Roadmap

During the first project year effort will be devoted to updating and extending existing OLA(s) to adapt them to the EGI needs. In addition, with the goal of delivering a reliable infrastructure which offers a smooth user experience, additional OLA(s) will need to be defined covering NGI and EGI core services.

A questionnaire will be distributed to NGI operations managers asking to provide input on requirements and extensions of existing OLA(s), such as increase of minimum thresholds for availability/reliability and hardware resources – see the Appendix for the full text.

For example, existing availability and reliability thresholds will be increased for existing production NGIs thanks to the growing maturity of middleware and operations, and these can be differentiated for new sites and NGIs during the initial learning curve.

The 1st line support OLA will also need to be adjusted to better reflect the EGI 1st and 2nd line support structures, to define minimum quality parameters for all existing support units, and to take into account for user support activities carried out by the individual NGIs with different commitments.

5.1. EXTENSIONS OF EXISTING OLAS

5.1.1. Tuning of thresholds

A recommendation has been made that the minimum availability and reliability monthly limits are slightly increased to 80% and 85% respectively. This increase reflects the progress that has been made in infrastructure performance throughout the EGEE project series.

For new NGIs and sites different thresholds might be applied for a grace period of six months (see Section 5.1.3 for more details).

5.1.2. Cases for site suspension

Current procedures require a site to be suspended if availability and reliability drop below 50% for three consecutive months. The Site-NGI OLA will be expanded to include a comprehensive list of cases for suspension such as:

- The site is suspended if affected by highly critical security issues (as assessed by the EGI CSIRT) and is failing to apply the required countermeasures within the required timeframe.
- The site managers do not handle operational tickets properly, for example they repeatedly fail to reply to trouble tickets in due time.
- A service running a retired middleware version is suspended if not upgraded after a grace period.

5.1.3. Differentiated Quality of Service

The current Site-NGI OLA defines standards minimum quality parameters that apply by default to all certified resource centres that are part of the production infrastructure. It is envisioned that in the future, minimum reliability and availability targets could be differentiated depending on the type of site and service operated, and on the quality parameters requested by the Virtual Organizations (VOs) supported by a given site.

- **Newly certified production sites.** For a limited grace period (for example six months) less stringent suspension procedures might be applied to services and sites that have been recently integrated into the production infrastructure while site managers get familiar with operational best practices and procedures.
- **User perception of availability and reliability.** Availability and reliability thresholds need to reflect not only basic service functionality – which can be tested through an operational VO such as “ops” – but also the user perception of the infrastructure. This implies the possibility to customize the availability and reliability of a site to include VO-specific test results into calculations through probes that better reflect the user workflows. These VO availability and reliability statistics can be then combined with the already existing monitoring VOs (ops, or the NGI monitoring VO) availability figures. To this end, it would be useful for example, if the operational tools could be instrumented to give the VO that possibility to declare the minimum quality parameters that resource providers need to provide. This information can be declared within the VO ID card, and can be extracted by the availability calculation engine. VO ID cards need to be periodically refreshed, and consequently the requested parameters might evolve over time. If a site supports different VOs each with an own minimum availability/reliability requirement, then the resource centre passes the availability/reliability test only if all requirements are satisfied (logical AND).
- **Differentiated sites.** It is expected that not all sites can offer the same performance given effort and experience constraints. These properties define the site profile, which might be combined with VO availability/reliability requirements in order to ensure that VOs are supported only by sites whose profile matches the VO minimum requirements.

- **NGI helpdesk response times.** NGIs will be requested to subscribe an OLA to commit to a minimum level of service requested by EGI to be part of the infrastructure. Nevertheless, not all NGIs are expected to provide the same amount of effort for user and operational support. Thus, response, solution and inactivity thresholds might be customized for each NGI according to its own needs, provided that the baseline EGI requirements are met.

5.1.4. New services

Availability and reliability of sites are currently measured against a minimum fixed set of grid services that are equally requested from all sites.

- Availability and reliability statistics need to include additional types of grid core services where available, such as top-level BDII, VOMS, WMS/LB etc.
- Novel types of services may be provided in the future as EGI is expected to integrate with novel Distributed Computing Infrastructures (DCIs). For this reason, the set of baselines services tested for availability and reliability statistics should be extensible, and generally speaking customizable according to the resources and services provided by a grid site.
- Network availability. The current Site-NGI OLA requests the quality of network connectivity to be sufficient to provide reliable access to the grid services exposed, but the network quality is not either qualitatively or quantitatively defined. Network bandwidth and average connectivity uptime can be included in the Site-NGI OLA. The Downcollector tool can be used for network downtime monitoring [R 16].

5.2. NEW OLA(S)

EGI is gradually evolving into a service-oriented infrastructure where part of the services are centrally provided, whilst others are distributed and operated under the technical responsibility of individual Virtual Research Communities and NGIs. The ultimate quality of service perceived by the grid user depends on the overall quality provided by the combination of all services, and it is therefore important that different service providers commit to a minimum set of requirements.

This section provides examples of new OLAs that could be adopted in EGI. New OLAs will be developed with the collaboration and consensus of all relevant parties. A questionnaire will be distributed to NGIs during the second quarter of the project to collect input. Such results will be processed and discussed in the framework of an Operational Level Agreement workshop organized during the first EGI Technical Forum in September 2010.

5.2.1. NGI-EGI OLA

NGI and EGI mutual commitments will be reflected in a NGI-EGI OLA that defines the set of services provided and the minimum requirements to be honoured by both parties. The OLA will define the services that the NGI provides to the EGI user community, and the global services that EGI provides to the NGI to ensure its seamless integration into the pan-European grid infrastructure.

5.2.1.1. Part 1: NGI to EGI

This part will define the services the NGI needs to offer in order to be part of a pan-European Grid infrastructure. The list includes examples of both operational services for site managers and

users, and grid core middleware services. Generally speaking for each category, services can vary over time, so the OLA needs to be flexible enough to include only those that are supported

- **Operational services**
 - NGI monitoring infrastructure
 - NGI Accounting infrastructure
 - NGI helpdesk
 - Other regionalized tools when available (GOCDB, dashboards etc.)
 - Other services that will be identified in the process.
- **Core services**
 - WMS/LB
 - Central LFC
 - Top-level BDII
 - VOMS
 - FTS
 - other
- **Response times**
 - of the NGI CSIRT
 - of the NGI operations staff on duty in case of urgent tickets to be addressed, requests for upgrade and suspension, etc.

5.2.1.2. Part 2: EGI to NGI

EGI will offer central global services to all NGIs that are essential for interoperation and seamless integration of different infrastructure. All these services and the related quality can be quantitatively defined through OLA(s).

- **Central operational services**
 - Central monitoring infrastructure
 - Central accounting infrastructure (portal and databases)
 - GGUS
 - Dashboard and operations portal
 - GOCDB
 - MyEGI portal
 - Messaging broker network
 - Metrics portal, etc.
- **Core services**
 - VOMS for dteam and other VOs
 - UMD repositories

- Web services and other support tools, etc.

5.2.2. Site-VO OLA

OLAs can be established directly between resource provides and the VOs supported.

For example, the site-VO OLA can specify VO-specific minimum availability/reliability thresholds and quality parameters of VO-specific services. In this case, monthly site availability/reliability statistics will need to be computed by taking into account not only statistics for the VO “ops” but also the other VO supported. A VO may also wish to define additional metrics associated to own requirements which can be monitored through a VO dedicated infrastructure.

5.2.3. EGI-VRC OLA

EGI and specific NGIs may be willing to provide services dedicated and/or specific to Virtual Research Communities. Consequently, OLA(s) can be defined to specific also the quality parameters applicable to user community technical services.

- User community services
- application database
- VO dashboards (where applicable) and scientific gateways
- VO management tools and VO validation
- VO-specific support services
- training, etc.

5.2.4. Tool extensions

Availability and reliability calculation tools are currently focused on grid common services whose functionality is tested through the VO “ops”.

The extension of existing OLAs and the introduction of new OLAs requires significant extensions to the current metric report generators. Depending on the service, metrics can be measured trough external probes or by gathering internal status information. For example, while functionality can be monitored by issuing external probes, response time to tickets requires the capability in the helpdesk to gather internal status parameters.

Generally speaking, monitoring of operational tools will require specialized new sensors in addition to the existing grid middleware probes. The Nagios system is sufficiently flexible to allow for such specialized probes to be developed and integrated. The availability calculation engine will need to take these new results into account in for its calculations.

6. Conclusions

OLAs are needed in EGI in order to provide a set of high-quality services to end-users by defining the service quality levels and the responsibilities of the individual operational units.

EGI relies on the experience gathered during the EGEE project series and will adapt and extend the set of OLAs used to meet the requirements of a pan-European infrastructure based on the



composition of diverse Grid national infrastructure, different interoperable tools and different harmonized middleware stacks.

Evaluation of the OLA metrics will be done on a monthly basis in order to ensure accuracy, progress will be monitored and procedures to undertake corrective actions will be needed.

New OLAs will be defined in order to cover the interactions between NGI and EGI, and EGI and the supported Virtual Research Communities. NGI experience and feedback will be gathered through a questionnaire during the second quarter of the project as a starting point of this new operational development effort.

7. Annex A: NGI questionnaire

This section illustrates the OLA questionnaire distributed to NGI operations managers in July 2010.

7.1. OLA STATUS

1. Number of certified sites in the NGI
2. Number of sites that have already signed an OLA or comparable document
3. In case of a comparable document being used, describe deviations from the metrics used in the original EGI OLA document.
4. What is the main obstacle to the adoption of the OLA by all sites?
5. Which are the main considerations / objections of sites to the OLA?
6. Describe any modifications that you would consider to the OLA metrics definitions?
7. Are there any metrics that should be added/removed from the OLA? Include a brief justification for your answer.

7.2. 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)
9. What kind of rewards/penalties for sites would you consider for over/underachieving sites?
10. Do you find the current system for providing justifications for A/R failures adequate? If not why? What else would you use?
11. Do the justifications in general adequately describe the incident, main cause and the recovery strategy used?

7.3. MONITORING TOOLS

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

7.4. 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?
17. How do you (or would you) manage OLA(s) in your NGI?



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

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?

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

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