Subject: **ANSWER TO EAC FEEDBACK**

Author(s): **PEB members**

Distribution **PEB**

General comments

EAC Comment 1

- The EAC is very pleased by the momentum of work that has been achieved since the 1st conference at Cork
- There are both good management practices and strong motivation of the participants to match the project objectives
- A sustainable scenario for the next 4/8 years concerning the EGEE Grid production infrastructure must be developed

EGEE Answer 1

A pattern representing a possible solution is emerging in the form of the closer link between eIRG and ESFRI. This automatically brings in long-term "roadmapping" aspect to the work of the eIRG. The EGEE project has been very proactive in its support of the work of the eIRG, which would offer a natural channel for information exchange between stakeholders developing long-term vision and policies and the EGEE project.

EAC Comment 2

EGEE should not solve problems in an isolated manner that others have to solve

- Security, network, …

EGEE Answer 2

1. EGEE COLLABORATIONS

EGEE is engaged in collaborative activities with other projects and initiatives on several fronts:

- Providing services to user communities together with other initiatives
- Using the same software components – either developed by collaborators or third parties. In case of third party software this means that both parties are aware of sharing.
- Information exchange and work towards interoperability in software development
- Participation in standardisation forums

These activities span different application domains and geographical areas (Northern and Latin America, Asia-Pacific region).

1.1. ACTIVITIES IN DETAIL

1.1.1. Service provision

EGEE has taken a contact with GN2/SA3 for the Network Operational Interface concerning the Trouble Ticket System which is being developed in GN2 by the PERT Team. The aim is to work with GN2 and FZK (GGUS) and define network interface and procedures between EGEE and GN2 (including also the NRENs). In addition to the common SLA definition work[1] with EGEE, the Technical Network Liaison Committee (TNLC) has been established by EGEE and GEANT/NRENs dealing with practical issues related to computer networking.

[1] See section on "Information exchange, interoperability".

In the field of security, the following areas are common activities with various US grid projects:

- Incident response plans
- Joint Security Policy Group

EGEE has also proposed common security service challenges - starting with walkthroughs of incident response procedures.
Furthermore, interoperation between storage systems has been demonstrated and is in production between the EGEE/LCG and various US groups. This work has also standardisation and joint development aspects (see SRM under standardisation).

1.1.2. Sharing software components

Grid3 project is interested in deploying the BDII information system. EGEE would like to bring the Grid3 information system and the LCG-2 IS together using the BDII and either filters to adapt the different schema interpretations or by ensuring that the GLUE schema (which is used by both Grid3 and LCG-2) evolves to include the requirements of both Grids.
EGEE and the Open Science Grid use the same basic authentication and authorisation infrastructure, development of which has been seen as a potential area of collaboration.

1.1.3. Information exchange, interoperability work

The work on Service Level Agreements (SLA) in networking involved a close relationship with the GN2 because it is necessary to have common criterions in SLA processing and their monitoring. EGEE JRA4 has also answered the GN1-JRA1 questionnaire regarding the networking requirements. EGEE and GN2 will have to agree on a common set of metrics to be provided, and on the definition of those metrics. GN2 will provide to the EGEE NPM prototype access to monitoring information through the NM-WG interface as soon as the interfaces is be defined and the measurement information collection mechanism is integrated within the GN2-JRA1 architecture.
EGEE and the Open Science Grid recognise that they have very similar problems and issues to address in the field of operations monitoring. Again, agreeing common schema to report the problems will allow viewing across the grids. The other issue would be to agree common metrics for operations performance and reliability. This could lead to development of common monitoring "dashboards" and/or common site validation tests (perhaps with an approach of a common sub-set of tests). Perhaps an even more fruitful area of collaboration would be common application validation suites for a site (if feasible).
EGEE and the Open Science Grid projects have also very similar problems to solve in accounting, although the mechanisms are different. Also in this case a strategy as mentioned above can be used - the Grid3 MonaLisa service could export directly into R-GMA allowing a top level query across infrastructures. Both projects already use the same (GGF-recommended) accounting schema.
Discussions are under way to agree job flow states for the EGEE and the Open Science Grid projects. The idea is to build a common tool that can be used in the same way to view job states in either infrastructure. Effort has also been identified to evolve the GLUE schema as a joint activity between Grid3 and EGEE.
Once the information systems are better aligned between the Open Science Grid and the EGEE, cross-grid job submission should be technically possible. However, there are several policy-level issues that need to be addressed before this is possible in practice.
Issues related to dealing with multiple VO's and/or generic VO's are recognised as a problem and could be a point of collaboration between OSG and EGEE. In addition to VO support, the following have been discussed at recent EGEE/LCG and Open Science Grid Operations workshops:

- Use of similar site service level definition (site charter in Grid3)

- Agreement on providing operations coverage across infrastructures around the clock, using local time differences between Europe and North-America?
- Agreement on tools on reporting and tracking problems between grids (will be needed if we have cross-grid jobs)?

Collaborating on the cookbooks and playbooks that both parties need was also discussed.

1.1.4. Standardisation

Concerning networking, EGEE JRA4 Network Performance Monitoring (NPM) aims to standardise interfaces allowing transparent access of network performance information from Measurement Frameworks (such as the GN2-JRA1 framework). EGEE expects these frameworks to be widely deployed throughout Europe (and beyond) to perform network tests and provide network performance information.
EGEE JRA4 NPM expects that these frameworks are able to provide performance results through the standardised GGF NM-WG Schema interface. They define that an NM-WG Report that provides the network measurement data can be provided through the submission of a NM-WG Request which gives information on the nodes, characteristic, time etc. that the user is interested in.
As mentioned in the context of interoperability and shared components, both EGEE and Open Science Grid use the same GGF recommended accounting schema.
In the storage systems, SRM is the de-facto standard that most of the production-oriented Grid projects are deploying. Both EGEE/LCG and U.S. groups work together within GGF and outside to develop this standard and its implementations.

Foreseen risks

EAC Comment 3

There is a high risk that the project may not meet its objective due to conflicting requirements and interests in the development of the gLite middleware

EGEE Answer 3

One of the primary problems with getting a consistent, global view of the requirements for EGEE software is that those requirements exist in many widely-dispersed documents. Early in the project, NA4 took the lead in using a requirements database to manage the NA4 requirements. Other activities have expressed an interest in using the tool, so this will become a project-wide database. This centralisation will help in forming a global view of the EGEE requirements.

Part of the PTF mandate is to review and prioritise those requirements. It has already begun the process. Unfortunately the existing requirements vary wildly in their usefulness to the developers and testers. The last PTF meeting focused on how to make the requirements as concrete and useful as possible. We reached a consensus on this and the existing requirements will be updated and the new requirements from other activities will be added early in the New Year.

Nevertheless, we need to be able to stand back between releases and bring all the inputs together (e.g. applications, deployment, development) and prioritise the next set of features accordingly.

NOTE: This text will be updated following the conclusions of the Management Task Force (see answer 4).

EAC Comment 4

The project is facing a difficulty in the development of gLite with two possible scenarios

- Focus JRA1 integration and testing on AliEn components
  - High-energy physics application will take benefit of such a scenario
- Continue delivery to pre-production service as planned
  - Most of the applications will benefit of such a scenario

Such situation must be addressed urgently by the Project Director having in mind the objective of the project

- "Enabling Grids for e-Science in Europe"
- We recommend thus to follow the second scenario

EGEE Answer 4

The Project Management Board unanimously supported the plan to adhere to the project work-plan (Annex 1) and ensure a release of gLite is ready for deployment in March 2005.

ALL effort (funded or unfunded, full-time or part-time) in JRA1 will be concentrated on bringing a selected set of high priority components to production-ready status. This process is currently taking place.

Any groups that wish to take earlier versions of gLite are welcome to do so but the support of these deployments is not the responsibility of JRA1.

NOTE: This text will be updated following the conclusions of the Management Task Force.

Middleware

EAC Comment 5

Need better support for the management of licensed software

EGEE Answer 5

It has been widely recognised that a Grid infrastructure such as EGEE should be capable of deploying applications based upon or containing proprietary software. A two-prong approach is being followed to address this issue. On the "experimental" side, two new applications deployed on EGEE fall in that category: EGEODE, containing a CGG-owned seismologic code and Computational Chemistry where the proprietary GEMS code is used. The various EGEE teams in charge of application deployments (mainly NA4 generic team and SA1 team) will work in close partnership with CGG to understand all the consequences of this situation. The results will be fed back to EGEE management and EGEE Industry Forum. On a more "theoretical" side, EGEE Industry Forum is now launching two working groups, one devoted to technological issues related to industrial usage of the Grid, chaired by Ian Guerin of IBM, the other devoted to business issues and economical models, chaired by M. Benard from HP. These working groups will make their first report in the next EGEE conference in Athens, which, together with the experience gained from the EGEODE case, will allow EGEE to define a solid policy on this important issue.

EAC Comment 6

The main requirement for a middleware is its robustness and stability

EGEE Answer 6

We need to focus on key essential features first and get them right (part of plan above). Avoid spreading our resources to thin on too many less important features. Need to make sure enough resources are made available to support testing and deployment efforts.

Testing teams (JRA1/Testing, NA4/Testing, NA4/ARDA and LCG-2/Certification) working together to thoroughly test the released gLite components/services. We also need to address configuration errors, which is to be addressed by the configuration strategy JRA1 is pursuing (How much of that will be in RC1?).

The Pre-Production Service is also an important tool to assess deployment issues, on a small scale.

We also need to involve more the Industry Forum, the eIRG and UK NGS.

EAC Comment 7

A common security infrastructure with other EU projects must be established

EGEE Answer 7

See point above on "General comments".

EGEE need to take, even further, a leading role on this topic. This also means promoting our already leading work better.

Standardising on VOMS and TLS is the current approach to establishing a solid and effective Grid security infrastructure for Europe, which can be suitable for the other projects on the timelines of the first EGEE project, including a common Authentication infrastructure (i.e. EUGridPMA).

In the medium term, we need to automate VO management in a robust and secure manner. The pseudonymity feature is to be provided though an external trusted third party that will have a role similar to that of a trusted CA and in the short term a client based encryption.
See further details from Ake in attached document:
http://agenda.cern.ch/askArchive.php?base=agenda&categ=a044542&id=a044542%2Fmoreinfo%2FEGEE_Security_Overview.ppt.

In the longer term, we should work with others (e.g. Grid projects, GGF, extended MWSG, TF-EMC2) to reach a more flexible authorization system that caters for other needs as well (e.g. digital libraries, student courses, etc) being in line with what's described in the eIRG white paper (v2.0). EGEE is also preparing a reach-out statement to all the EU projects related to security issues, in order to ensure that the practical experiences from the running of the Production Service are disseminated in the developer community. At the moment this action is pending the concretisation of the overall concertation plan of the EU projects.

In terms of implementation, MLS is likely to be a longer term solution, more adapted to interoperability with other grid infrastructures, through already established standards, such as WS-Security. Here too, this work will also require collaboration with other projects such as OSG, OMII, etc.

Support, training & Documentation

EAC Comment 8

- The project is producing a substantial amount of documentation and they must be kept updated
- The project involves a lot of participants. A "who knows what" register could help
- We recommend to build up a documentation tree for the VOs/users especially for new users
- Specific effort to be made to improve the documentation allowing autonomous installation of the gLite middleware

EGEE Answer 8

The EAC have highlighted an important challenge for EGEE. The requirement to maintain and develop a complete, coherent, up-to-date and consistent collection of documents describing EGEE is well recognised. It is particularly important that this collection is easily accessed by all varieties of new users and new site administrators. It should be well structured and well oriented to answering users' and developers' questions. Given the high accession rate for EGEE documents and their orientation towards internal EGEE goals, it is difficult to meet this information requirement.

A group, the "User Information Group (UIG)" with representatives from SA1, NA2, NA3 and NA4, has been set up to extract and develop a collection of documents that meet the requirements for new users and new VOs joining EGEE. This will, we hope, at least provide material in the appropriate form accessible via the main EGEE web site. It will also be made consistent with the currently available training materials repository. This will form the top level of the document tree, include access via FAQs and provide a focus for a search facility. In the longer term we hope to work with colleagues in the DILIGENT project to develop the information base and its supported management procedures more fully.

The collation of a directory of expertise within EGEE would be valuable. At present resources have not been allocated to collect and standardise the information – this should be addressed after the review perhaps by a survey with a categorisation of relevant knowledge devised by the UIG, so that it is user oriented.

Applications & Demonstrations

EAC Comment 9

Play an important role to demonstrate the capability of the underlying Grid Middleware and not only the applications themselves

- High responsibility of the persons who make the demonstrations

Demonstrations should be prepared with a common template

- Follow a unified approach

Applications must identify the benefits of a grid infrastructure compared to classical ones

- What have been made that was not allowed using classical computing infrastructure

The EAC found that gPTM3D was the most successful demonstration

EGEE Answer 9

In preparation for the review, we have put in place a comprehensive rehearsal programme, in order to make sure that the different elements of the demos and their accompanied presentations are of a high quality and address the right issues about EGEE. In preparation for these rehearsals, we have provided all presenters with a Power Point presentation template, including important slides to be filled in like the benefits of grids.

3 application demos, including gPTM3D, and the SA1 operations tools demo have been selected and will be presented during the review.