

Transitioning to a sustainable European Grid Infrastructure

July 3rd 2009

1. PROPOSAL OVERVIEW

Following a decade of successful research in Grid technologies, starting with experimental services in the EDG project and continuing with production-quality services for multiple user communities in the EGEE series of pioneering projects, the European e-infrastructure community through the National Grid Initiatives (the scientists that use it, the sites that provide the infrastructure and the developers that create the enabling middleware) are ready to consolidate their experiences and deliver a sustainable high quality e-infrastructure for the future to enable the European research community as described in the EGI_DS Blueprint.

This three year activity will allow the community to finish its transition from previous projects and consolidate and rationalise its activity into a steady sustainable state to provide:

- A governance structure that represents the major stakeholders involved in delivering a federated European e-infrastructure built from national resources and the research communities willing to use it.
- A scalable user support model that allows international research communities within Europe and their worldwide collaborators to use the resources within the e-infrastructure that they are allowed access to.
- Interoperation and interoperability with other e-infrastructures within Europe and with those globally required by the European research community.

EGI is a collaboration of National Grid Initiatives (NGIs) working together with a coordinating body (EGI.eu) to deliver a service to their own user communities and their collaborators. EGI will deliver consistent, high quality service and support through its own operations infrastructure and through the NGIs that interface with this European infrastructure. This investment in NGIs will be coupled with a management structure that will focus on, and deliver, a high quality service to the user community.

For EGI to succeed it will need to establish relationships with other EU projects and structures. Amongst these are the domain specific Specialised Support Centres (SSC) that provide the vital bridge between the application communities and the e-infrastructure. Relationships will also be established with software providers to ensure a consistent supply of high quality middleware releases that meet the needs of the user community in terms of stability and functionality.

The heavy multi-national users of federated distributed computing infrastructures expect certain capabilities beyond a set of reliable site-based services discovered through a central information service. They expect higher-level services that enable their access to site services, such as workload management services, file transfer services, meta-data catalogues, replica services, etc. These higher-level and the site services are critical to these user communities and the rapid investigation and resolution of critical issues found during their deployment and in production use within EGI needs to be assured.

2. PROPOSAL STRUCTURE

This section describes the structure of the project (work packages) and the tasks within each work package. Where available and relevant the role/task descriptions (e.g. D-E-1, O-E-1, etc) from EGI_DS D3.2 are provided (see section 9, Page 59).

This document should be read with the following considerations:

- These work packages and their associated tasks are indicative of the work to be undertaken in this area. At this stage of the proposal all of the work packages **will** need refinement following feedback from the community.
- Many of the tasks within the work packages are drawn from the description from the EGI_DS Blueprint.
- The term NGI refers to the federated entities within the EGI structure. It is recognised that some NGIs may choose to engage in their International Tasks on a regional basis, or to outsource these tasks to another body.
- In the context of providing an EGI Global Task the NGI could also be an EIROForum organisation, such as CERN.

WP 1: Project Management (NA)

1. Consortium Management (D-E-1, D-E-2): The support of the Director in their interaction with external bodies and with the EGI project collaboration.
2. Consortium Administration (D-E-6): The administration of the EGI.eu and the EGI project collaboration.
3. EGI Council Support (D-E-2): The support of the EGI Council that has the NGI representatives and other stakeholders.

WP 2: External Relations (NA): EGI's non-technical relationships with external organisations.

Managerially this WP will probably report to the Director.

1. Dissemination (E-E-1, E-E-3): Work within the project will be disseminated internally and to external media partners. It also coordinates dissemination activities at events attended by the dissemination contacts in the NGIs/SSCs (U-N-2 & U-S-2 – not funded through this project). It also collects results from these NGI/SSC representatives for internal and external dissemination.
2. Policy and Standards (E-E-4): EGI will be main high throughput distributed computing infrastructure in Europe. As such EGI will be influential and will need to influence the policy agenda in Europe and worldwide and engage in developing this policy alongside other European e-infrastructures such as DEISA and PRACE, and other international e-infrastructures such as TeraGrid and Open Science Grid. Likewise, the standards adoption and impact within the e-infrastructure undertaken by the software providers will need to be analysed.

3. User Community Relations (U-E-12 & E-E-3): The relationship of EGI with its User Community (including related projects within Europe and worldwide) is coordinated here. A class of key related project(s) are the SSCs whose collaboration with EGI is critical that the relationship is defined through the UFSC – so the enablement of this activity is explicitly supported. This task also includes the support and management of two large community driven events a year (a European e-Infrastructure Conference and a European Distributed Computing User Forum).

WP 3: e-Infrastructure Operations (SA)

1. WP Mgmt: Management of the whole activity is undertaken by the Central Operations Officer (COO) (D-E-4) who will be assisted by the person driving the grid operations and providing oversight of the e-Infrastructure (O-E-5) and the longer term development of the operations infrastructure through the gathering of requirements for user support tools and processes (O-E-8). The management of the staff based in the NGIs that are funded under task 2 will also fall under this task.
2. The NGI International Tasks: The funding for these tasks (O-N-*, tasks from EGI_DS D3.2) will be placed within the NGIs that are part of EGI with well defined SLAs. It is expected that the staff behind these support functions are not working solely on the specified functions (*-N-*) and therefore the contributions provided by EGI are a contribution to a fraction of their time in return through the work done through the agreed SLA. Consistent failure to meet the SLA could result in a proportional reduction in the funding for that function.

The nominated effort supporting the teams behind these NGI International Tasks is solely for generic infrastructure issues coming from the international collaborations being supported by the NGIs. It is NOT for domain specific application support, support of particular needs of a community, support of national projects, etc.

3. Operations Coordination: This task provides the central driving force behind the operations function within EGI. It includes triage of incoming tickets (O-E-7), coordination of middleware rollouts and the coordination of interactions between the software providers and NGIs as part of a staged rollout and pilot services (O-E-9), coordination of resource allocation and brokering support for the VOs (O-E-10), coordination of interoperations between NGIs and other grids (O-E-11), coordination of network support (O-E-12) and the coordination of best practice definition, operations procedures and requirements (O-E-13).
4. Security Coordination: Covering the development and maintenance of security policy (O-E-15) and the coordination of security and incident response (O-E-16).
5. Operations Services: This task includes operation of the grid topology and configuration repositories (O-E-1), accounting repositories for the international VOs (O-E-2), repository storing monitoring and performance data and related information (O-E-3), a Grid Operations Portal (O-E-4), a operation of a ticketing system and document repository for central and regional grid support (O-E-6), and the operation of a production grid core software services (e.g. catch-all VOMS, CAs, etc.) (O-E-14). These services will be undertaken by a small number of partners with well defined SLAs.

WP 4: Operational Tools Development (JRA)

1. WP Mgmt: Coordination and development of operational tools (O-E-17).
2. Maintenance and development of the operational tools: This is a continuous background activity that takes place during the project. The effort required by the teams providing the EGI Global tasks need to be defined depending on the tools and the partners selected through the bidding process.
3. Regionalisation of Operational Tools: This is a short term task (12-18 months) covering the regionalisation of the operation tools – which will not be fully complete during EGEE. As a result of the EGI Global Task bidding process we will know which tools will need to be further regionalised. Input will also be needed from the community as to which tools need to be developed beyond a regional level to a national level.
4. Automation of Operational Tools: This will be a continuous task during the project which will focus development effort on automating tools and processes that currently require a high manual input from NGI resources. The key motivation being to invest in development effort now to reduce the running costs of the infrastructure in the future, i.e. balancing the cost now against the future potential savings. The scope of this task will need to be coordinated with other work that may be taking place within existing projects or newly proposed projects. The work will focus on practical enhancement to the tools that reduces ‘human’ workload – a more research oriented agenda looking to exploit autonomic operation models is not seen as in scope.

NOTE: Tasks 3 and 4 of this work package will require an assessment from the current community as to the state of regionalisation and automation of the operations tools when EGEE ends in April 2010.

WP 5: User Community Support (SA)

1. WP Mgmt: User Community Officer (D-E-5)
2. User Coordination: This task includes coordination of the SSC activities (U-E-10), coordination of the front desk provided by the support services (U-E-13) and coordination of the training effort (U-E-15).
3. User Support: Activity covering the support of new & small communities (U-E-11) and the documentation of User & SSC interfaces within EGI (U-E-14).

NOTE: Tasks 2 and 3 of this work package need to be further developed with the SSCs to define the interaction between these tasks and the functions of the SSC.

WP 6: Software Infrastructure Coordination (SA)

1. WP Mgmt: This activity is managed by the Central Technology Officer (CTO) which is the role (D-E-3) and this task also includes the management of software providers (M-E-5).
2. Provision of Technical Services: This task provides technical services to EGI and the software providers that EGI works with. One aspect of this work is a repository of components (M-E-4) used as the ‘central’ repository of certified binary and source components available for deployment on the NGIs. This repository could also be a source of uncertified components contributed by the community (along the lines of the RESPECT project) using a process similar to RESPECT driven through an SSC oriented process (e.g. UFSC). An additional activity is to provide and support process

tools (M-E-2) for the middleware unit and EGI. For the MU this could include an issue tracker, more generally for EGI this could include web sites, wikis, meeting agenda.

3. Criteria Definition and Verification: To support the inclusion of software into UMD, the software providers will have to produce software that will need to meet defined criteria (see process document) which will be verified by the MU (M-E-3). These criteria will be defined and developed, as long with related process document (M-E-1), in consultation with the EGI's software providers as represented in the MCB.

WP 7: Support and Maintenance of UMD for Heavy Multi-National Users of EGI (SA)

1. WP Mgmt: The activity will be managed on a daily basis by the MU (M-E-5) who will triage, assign and follow up reported middleware issues. Administratively this project will be managed by the CTO.

The middleware being used in the production infrastructure is a critical element in providing a quality production infrastructure. Critical issues are frequently found in the production use of the middleware and these need to be rapidly investigated by experts that are associated with the production service. These issues may be simple 'bugs' or relate to the way a user is using (or misusing) the software. Such 1st line middleware support mechanisms are vital to delivering a high quality production grade service yet are not systematically provided within EGEE. Instead, such support is found in some communities through their own experts where it is seen to be highly effective. This work package builds on this experience and makes it available to other important user communities.

A 1st line middleware support unit attached to the production infrastructure would work closely with the 2nd line support units with the individual middleware providers through coordination by the Middleware and Operations units within EGI.eu. The 1st line unit would be able to undertake preliminary debugging of issues, either identifying fixes for the simple issues or by narrowing down the area of investigation needed to be undertaken by the 2nd line units.

This work package does NOT remove the need for a broader middleware activity with (the 2nd line maintenance activity) that will be part of the work undertaken by the European Middleware Initiative (EMI) consisting of the gLite, UNICORE and ARC consortia and other software providers which will also focus on the harmonisation (standards definition and adoption), convergence (component combination), and development of new functionality needed by the EGI community to meet their long term needs.

While the teams in Tasks 2-7 may be in different locations each team will be co-located within a single partner (ideally within a single office!) to ensure minimal overhead and maximum effective in the middleware maintenance activities. These teams will work closely with the teams within EMI (and maybe collocated with them) contributing fixes to the production release branches of these components that can then be taken up the EMI developers in their next release.

2. Maintenance and Support of the Security Components in production use.
3. Maintenance and Support of the Compute Components in production use.

4. Maintenance and Support of the Data Components in production use.
5. Maintenance and Support of the Accounting Components in production use.
6. Maintenance and Support of the Information Components in production use.
7. Support for component build and integration.

WP 8: Services for Heavy Multi-National Users of EGI (SA)

1. WP Mgmt: Coordinate the operation of services for heavy multi-national users of the EGI and to develop this service provision in response to the needs of this community.
2. Service Hosting: Provision and operation by a small number of NGIs of Core Grid Services (O-N-8) explicitly needed to support this user community, but of potential benefit to other communities. These centres will be experts and provide an SLA around the hosting of services *such as* FTS, LFC, Hydra, AMGA and VO specific services. This list is an **example** of services which heavy multi-national users of EGI will be critically dependent on and their service provision by experts within some NGIs supporting these SSCs needs to be assured. The services and the hosting locations will be further discussed and the task developed with the relevant SSC communities.
3. VO Dashboard: VO Dashboards have been found to be very useful by large VOs to provide a VO view of the infrastructure for their community. This task includes the hosting of the service and the integration and development of VO specific tests, driven by the particular user community, necessary to verify the correct functioning of the infrastructure for their work. This will also draw on the generic service monitoring infrastructure and tests maintained by the NGIs.
4. Scientific Gateways: Provision of a generic science gateway infrastructure available to heavy users for customisation and extension. The task would provide the expertise to host the service and to deploy established components in response to the requirements from the community leaving the science communities (SSCs) free to concentrate on science and not service hosting.

3. TASK/ROLE REFERENCE

3.1 NGI INTERNATIONAL TASKS

NGI International Tasks are undertaken by an NGI to interface their national activities into the European e-Infrastructure. For further information on these tasks see EGI_DS D3.2.

O-N-1	Operation of the NGI Grid topology and configuration repository - <i>necessary</i>
O-N-2	Operation of the NGI accounting repository - <i>necessary</i>
O-N-3	Operation of repositories storing monitoring and performance data, and other related information – <i>necessary</i>
O-N-4	Operation of the NGI Operations Portal – <i>necessary</i>
O-N-5	NGI e-Infrastructure oversight (monitoring of status of services operated by sites, opening of tickets and their follow up for problem resolution), 1 st and 2 nd line support in case of operational problems, site suspension, reporting to EGI.org in case of middleware problems and general operational issues, etc. – <i>necessary</i>
O-N-6	Operation of the NGI ticketing system, gathering of new requirements for support tools in the region – <i>necessary</i>
O-N-7	Regional helpdesk: support to users and site managers via a local/regional helpdesk and documentation - <i>necessary</i>
O-N-8	Operation of production Grid core services, catch-all services for international VOs, catch-all CA: running the required Grid services provided by the NGI, and services required by international VOs – <i>optional</i> ; availability of Certification Authority: to distribute X.509 certificates to users and servers in the region - <i>necessary</i>
O-N-9	Operations Coordination at the NGI level - <i>necessary</i> <ul style="list-style-type: none">a) Security and incident response coordination in the regionb) Roll out of middleware updates in the NGIc) Resource allocation in the NGId) Interoperation with national and regional Grids

3.2 EGI GLOBAL TASKS

The EGI Global Tasks are activities undertaken by an individual NGI on behalf of the EGI community because they have particular expertise or are able to offer particular service offering. These tasks have been defined in EGI_DS D3.2 and are included here for reference.

O-E-1	Operation of the grid topology and configuration repositories	1
O-E-2	Operation of accounting repositories for international VOs	1
O-E-3	Operation of the grid repositories storing monitoring and performance data, and other related information	2.5
O-E-4	Operation of the grid operations portals	0.5
O-E-5	Grid operation and oversight of the e-Infrastructure	1
O-E-6	Central and regional Grid support, operation of a ticketing system and of a document repository	2
O-E-7	Triage of incoming problems	2
O-E-8	Gathering of requirements for user support tools and processes	0.5
O-E-9	Coordination of middleware roll-out and deployment, middleware pilot and certification testbeds	1
O-E-10	Coordination of resource allocation and of brokering support for VOs from NGIs	0.5
O-E-11	Coordination of interoperations between NGIs and other grids	0.5
O-E-12	Coordination of network support	0.5
O-E-13	Coordination of definition of best practices, operations procedures, operations requirements	0.5
O-E-14	Operation of production grid core software services, catch-all services for international VOs, catch-all CA	1
O-E-15	Coordination of security policy development and maintenance	0.5
O-E-16	Coordination of security and incident response	1
O-E-17	Coordination of development and maintenance of operational tools	1

