

Greetings to the EGEE-III SA1 Activity managers, and the EGEE Community,

Recently a “*proposal for the centralized distribution of gLite client software to EGEE sites*” [CENTINST] was circulated to the resource centres participating in EGEE-III on behalf of the SA1 activity management, proposing the use of human and technical project resources for centralised distribution of grid middleware from a single source by the SA1 activity in EGEE-III, and requiring resources centres participating in the project to accommodate the use of this centrally installed middleware by end-users and user communities.

Considering that

- (i) The EGEE-III project has as one of its principal objectives “*to prepare the migration of the existing Grid from a project-based model to a sustainable federated infrastructure based on National Grid Initiatives*” [DoW];
- (ii) Both the project reviewers, the European Commission representatives, as well as the Project Management Board have endorsed this objective of a federated infrastructure;
- (iii) The European Grid Infrastructure (EGI) design study (EDI-DS) foresees the e-Infrastructure in Europe in 2010 and beyond to be based on a distributed and federative model;
- (iv) The EGI Blueprint (draft) [D44] explicitly states in the Statutes (section 7.1.1, article 1) that “*the sharing of tasks between the National Grid Initiatives and the EGI.org shall follow the ‘subsidiarity principle’...*”;
- (v) Any proposal to centralise tasks, including but not limited to the installation of grid middleware, that have up to now been effectively performed by (the majority) of NGIs, federations of NGIs, and individual resource centres, is clearly and obviously incompatible with the aims of the EGEE-III project and the future EGI e-Infrastructure;
- (vi) That the signees of this position statement endeavour to make the EGEE-III project successful in attaining all of its stated goals;
- (vii) That the signees are committed to creating a sustainable and persistent e-Infrastructure that is beneficial for both national, regional and multi-national communities of e-scientists;

the signees would like to point out the following procedural and policy issues regarding this proposal, and request that (i) work on this proposal be stopped and (ii) the EGEE-III SA1 Activity managers carefully review any future proposals with respect to the Objectives of the project and the aim of creating a long-term sustainable e-Infrastructure.

As is clear from the Description of Work (DoW) and the EGI-DS Blueprint draft, the upcoming e-Infrastructure in Europe is based on a federative approach, such that the structure is self-sufficient and sustainable. The devolution of tasks that can be done at the national or regional level is essential to attaining this sustainability, since direct funding for central activities will diminish, and in the future be at least partly drawn from contributions by the NGIs.

Being mentioned explicitly in the DoW, this overall direction of EGEE-III has been clearly communicated throughout the project, and we therefore assume that this position on subsidiarity, federalisation and devolution of tasks to NGIs and regional organisations (such as ROCs) is also a guiding principle for the SA1 activity managers.

The apparent endorsement by the SA1 activity managers (with “EGEE Operations Management” as a signee of the proposal) must therefore surely be an oversight on behalf of the activity managers. A proposal such as ‘centralised software installation’ is obviously contrary to the principle of federation and subsidiarity.

The adoption of this proposal would create a de-facto reality that would subsequently compel a future EGI.org to provide this ‘service’ as part of its operational tasks, and in the structure envisioned for the EGI the NGIs would be paying for this ‘service’: obviously this service has a cost, which we estimate to be between 1 and 2 full-time equivalent, that in one way or the other has to be paid. From which cost centre in EGI it will be paid will in reality be mostly irrelevant. It is unlikely that this ‘service’ will stop again once the EGEE-III project comes to a close, or otherwise it would be unreasonable to establish this ‘service’ if its foreseen life time is only 12 months.

It is especially for this reason that any such form of central installation by SA1 must be considered unacceptable *ab initio*, and using or proposing to use project resources in a way that is so obviously contrary to the aim of a sustainable federative e-Infrastructure might be considered irresponsible and should be prevented by SA1 activity management (and ultimately by the PMB).

At the same time, devolution of responsibility from ‘centrally managed activities’, such as infrastructure monitoring, incident follow-up and end-user support, is already taking shape, with federations and NGIs performing tasks that up till now were done centrally – a move that is likely to benefit not only the sustainability of the infrastructure, but at the same time improve quality, reliability and accuracy of monitoring and incident handling. To at the same time move software installation from a federative and distributed model to a single ‘central committee’ comes at a strange moment at the very least.

Of course there are also technical issues that make this impractical and detrimental to the users. A large number of resource centre representatives have clearly articulated other concrete and specific objections at this technical level.

There are very good reasons why resource centres test - and based on these tests possibly delay – the installation of ‘new’ gLite middleware. It is not uncommon that software provided is not yet fully

stabilised, or does not integrate with the local environment. For instance, updates to middleware and external dependencies may adversely impact middleware used by regional or national communities, or it adversely interacts with the site's operating environment. Users that will inadvertently happen to use the 'central middleware', or are otherwise 'enticed' to use such an installation, will experience severe disruptions that not only impact that own work, but will also significantly increase support load on the resource centres, ROCs and NGIs. It is for this reasons that many sites use a local certification process to evaluate consistency and prevent such incidents from affecting the scientists that rely on the e-Infrastructure to work and do their research. The confirmed disruptive effects to the local user community in the federation and region, and the adverse effect on the amount of support staff needed at the sites make it hard to see any technical benefit in this proposal for central software management by SA1 for this reason alone.

The proposed mechanism will necessarily rely on a simple file distribution mechanism and as such cannot do dependency checking. This means that the consistency of the installed software cannot be guaranteed, which is a step back from the current situation where automatic tools exists to ensure such consistency. Necessary site-specific configuration changes to the middleware also cannot be accurately taken into account, and will adversely affect usability and completeness.

Any incorporation of a mechanism to select centrally installed middleware in the gLite workload management system will entice users to select the use of such middleware. Since this middleware cannot be complete, and will lack information on local VOs and pre-empt local additions from taking effect, these VOs will experience defects and therefore request local support.

For larger installations scaling issues in the way the centrally installed software is made available to individual 'worker nodes' will arise. These issues are already present today for software installed in the VO-specific 'software areas' for large sites and large VOs, and since the proposed mechanism for central middleware installation is based on the same technology, hitting these scaling limits is practically guaranteed.

At the same time, end-users that would be enticed to use the 'central installation' will experience issues and incidents that would not have occurred if the site-provided and certified installation would have been used. However, one can safely assume that users and VOs will address the site with any incident and problem they encounter, irrespective of the middleware installation that was actually used. Problems will be reported to sites based on the 'central installation', over which the site has no control, and whose issues therefore by definition should never be directed to the site. Since users will not appropriately direct their complaints, this will incur an unnecessary and unacceptable load of the end-user support services of the site, NGI or region.

However, if SA1 were to establish a central 'help desk' for problems with the central installation, this would require human resource to staff it, thus creating yet another centralized 'service' for which sustainability and persistence beyond the project lifetime is to be ensured. As with the installation service itself, creating such a central service is in clear violation of the Objectives of EGEE-III and a

violation of the subsidiarity principle. Thus, a central 'helpdesk' by SA1 must not be provided and is unacceptable.

Of course, it would be entirely within the remit of SA1 to propose that a mechanism for coordinated installation be developed and deployed. In fact such a mechanism is already in place and in use, and user communities today use this mechanism to distribute their own additional software suites. We acknowledge that a mechanism to distribute also grid middleware may be of benefit to selected smaller sites that do not have a local user community and whose operating environment does not deviate significantly from the 'boilerplate systems' used as templates in the certification of the gLite software.

However, no body 'central' in SA1 must take on the role of distributing software, or request authorization to do so from individual resource centres. Such a distribution to selected sites must be initiated from the NGIs, regions and federations, also because they will be better placed to assess the effects on specific sites in their country or region. Doing the actual software distribution 'centrally' from a single point in SA1 is fully outside the scope of EGEE.

Alternatively, an improved release mechanism that would allow sites to automatically install updates (technically implemented through auto updates using 'yum' or 'apt', such as is common for the GNU/Linux operating systems themselves) could be pursued. Enabling auto-updates would require that the distribution quality is improved (it is currently discouraged by the SA3 activity to enable auto-updating, as was recently re-asserted when an incomplete gLite 3.1 updated was posted in the download repositories [LCG-ROLLOUT-20080819]), but if this done correctly it would address the issue which 'central installation' is attempting to solve in a consistent and industry-standard way.

In summary, based on both the procedural and technical reasons stated above, the SA1 activity and the SA1 activity managers must not pursue the proposal for centralised software installation, and SA1 activity management must keep the principle of subsidiarity, the federated nature of the European e-Infrastructure, and the overall goal of sustainability clearly in mind when formulating any future proposals and policy.

With Kind Regards,

The Dutch Resources Centres participating in EGEE: Nikhef, SARA, RUG-CIT
The BiGGrid project Executive (Dutch NGI)
The Nordic Data Grid Facility (NDGF)
KTU Elektronikos Projektavimo Laboratorija, Lithuania (Site KTU-ELEN-LCG2)

cc: the Benelux Federation Representative in the EGEE-III PMB

References

- [CENTINST] Nick Thackray et al. "*A proposal for the centralized distribution of gLite client software to EGEE sites*", EGEE Broadcast dd August 28, 2008.
- [DOW] The EGEE-III Consortium "*Annex I - "Description of Work"*" to grant agreement 222667, May 15, 2008-08-31
- [D44] The EGI-DS Consortium et al. "*Draft EGI Blueprint Proposal: EU Deliverable D4.4*", June 23, 2008
- [LCG-ROLLOUT-20080819]
LCG-ROLLOUT mailing list, response to "*Silent gLite 3.1 update?*", August 19, 2008.