

Summary of concerns and recommendations of the Internal non-Apps Review 2003

Comments: Ian in blue, Frédéric in red, Les in green

Formatted: Font color: Green

These responses have been prepared for discussion at the PEB on 2 March. The purpose is to record the main recommendations of the internal (non-applications) review, and state the reaction of the project – if any.

Formatted: Font: Not Bold

Grid Deployment

- We recommend that a specific small “validation testbed” is created (at CERN) with the specific purpose to let the experiments test their software before official LCG releases
 - This was done in December (and in fact had already been started at the time of the review), and is now fully exploited by the experiments preparing for the data challenges.
- The Tier-1 centres have had many problems with the installation of LCG-1. This shows that there are issues related to the installation and configuration that were not caught during the test phase. The LCG area/task should pay closer attention to this issue.
 - The stability of the configurations is addressed in LCG-2 by starting with a set of core sites and ensuring that they have sufficient dedicated effort available to manage the service and its configuration. Sites will be integrated into the core after strict verification using a test suite which has been developed. In addition, sites will be pro-actively removed from the production core if problems are seen that cause problems to the system. We now have tools within the information system to do this simply which did not exist at the time of the review.
 - We have simplified the installation of the middleware on the worker nodes for LCG-2. We will continue to address and simplify the installation of the service nodes, reducing the dependencies on complex tools.
- We recommend that closer links between the relevant actors be put in place. Current GDB meetings do not include “the troops” – perhaps a regular (monthly/bi-monthly) meeting of a technical nature would help.
 - We now have a weekly Grid Deployment Area coordination meeting to which all stakeholders (users, sites, deployment team) are welcome. This meeting focuses on technical issues, referring policy matters to PEB or GDB as appropriate.
 - There is a weekly phone conference between the site system managers and the deployment team to specifically address deployment and installation issues.
- Particular attention should be paid to the issue of newcomers (and synchronizing them to the old and established practices).
 - Most Tier 1 sites are now integrated into LCG. Newcomers should in future be supported by their Tier 1’s exactly for this reason. In addition there are full installation instructions and release notes for each distribution.
- What is needed is a clear strategy towards full tests of the computing models before the LCG TDR. Identified as a “global issue”.

- The aim in 2004 is to test baseline computing models at the Tier 0-1 level, perhaps including large Tier 2s. These models will probably not be the official computing models of the experiments, even in this limited domain – rather they will be the base fallback models for reconstruction, reprocessing and batch analysis of the ESD.
- More advanced computing models are not scheduled to be defined before the end of 2004. Are the experiments on time for this? (question from the POB).
- The ARDA project will focus on testing possible advanced analysis scenarios. It is too soon to see how far this will evolve into a solid strategy for the experiment computing models.

Formatted: Bullets and Numbering

Formatted: Font color: Green

Middleware

- While the M/W is not under the exclusive control of the LCG project, its milestones are very important and need to be included in the project overview.
 - The milestones are being included in the project overview (see last quarterly report).
- ARDA planning should be established by end 2003, involving both the experiments and EGEE m/w experts, as well as AliEn, NorduGrid and US m/w experts.
 - An ARDA project has been formally proposed mid-February 2004. The EGEE middleware core team currently involves members from AliEn, EDG, EGEE and VDT. Other contributions are expected and being worked at.
- The six-month timescale for the ARDA prototype should be negotiated with EGEE and the experiments: Real point: to have a new release for users before end 2004
 - The timescales are currently to provide a first version of a prototype in Spring 2004, followed-up by rapid upgrades cycles. This will be the base for EGEE middleware.
- Federated/multiple grids - First priority should be to show that a single Grid can achieve real production quality - this is the LCG.
 - Not sure I need to add something here.
- A fallback solution for Grid m/w is very important - especially if LCG-2 evolution does not deliver production-quality m/w in time for the experiment C-TDRs
 - This is tricky, I do not think this comes from Middleware only, more realistically with a combination of experiments and LCG ad-hoc solutions?.
 - I think that LCG-2 is the fallback solution – going further back than this does not address data and storage management. The fallback is rather to restrict the scale of the grid – e.g. Tier 1 only.

Formatted: Font color: Green

Fabric

- No major concerns.

Management

- Lacking authorization for phase 2 of LCG, the long-term support of software packages developed in EDG and LCG is a concern. This needs to be addressed in 2004, well before the end of phase 1

Les Robertson
24 February 2004

- An agreement is being worked out now for medium term support of the EDG and VDT software components in LCG-2, with source code level fixes possible by the LCG Grid Deployment team at CERN (EDG) and the VDT team at Wisconsin (VDT, Globus), limiting the need for recourse to experts. Support after 2005 is not clear yet – possible solution is continuing GD/VDT support at a cost of a few FTEs – feasibility will only become clear with experience later this year. However - current strategy assumes that this middleware would be replaced by EGEE package before end 2005.
- Long term support for EGEE tools – not clear. This must be a consideration of the EGEE middleware development team, but the timescale for resolving this is clearly beyond end 2004. This will however be a criterion for the decision to replace LCG-2 with EGEE tools for mission critical (Tier 0/1) applications.
- Applications software – Review of the AA work plans is under way at present. Following this we will review the long term requirements and available funding. This is part of the full Phase 2 planning that is scheduled to be completed by the summer of 2004.
- The relationships between LCG-2 and the post-prototype ARDA implementation should be stated more clearly
 - I think that this is now being done – LCG-2 middleware will be solidly supported by a team independent of the EGEE middleware developers, in parallel with development and deployment of EGEE middleware. This will continue until the LCG-2 middleware is replaced by new tools.
- The manpower situation seems to be almost OK for the next year, but in the longer term there are many problems, which may lead to an untenable situation.
 - This is being addressed as part of the general Phase 2 planning that is scheduled to be completed by the summer.
- The Computing MOUs, synchronizing with the Computing RRB, and the overall plan, etc is a huge issue.
 - Indeed it is. This has been taken up as a priority by the CSO.

Formatted: Font color: Green

Formatted: Font color: Green

Formatted: Font color: Green

Formatted: Font color: Green