Status of openlab-I and Plans openlab-II

Platform **C**ompetence **C**entre
GRID **I**nteroperability **C**entre

Version 2
22 October 2005
François Fluckiger, Sverre Jarp

This document summarises what the CERN openlab industrial partnership has achieved to date, and current plans for new projects in the period 2006-08.

Short review of “openlab I”

Principal mission

The CERN openlab for Datagrid applications was initiated during 2001-2002 and formally launched in January 2003. In order to evaluate and integrate cutting-edge technologies focused on the needs of the LHC Computing Grid (LCG), leading IT companies with a solid track record of delivering advanced technological solutions were invited to join. The guiding principles for joining are available as a Web document. Since 2003 CERN openlab has involved CERN working in collaboration with Enterasys, HP, IBM, INTEL, and Oracle on a three-year project called the “CERN opencluster”.

Funding model

The activities of openlab are funded by contributions from each of the partners, which can be in-kind (equipment, manpower) and which amount to at least 2.4MCHF per partner over three years. There is also a contributor status for smaller firms at a contribution level of 240kCHF for one year. The company Voltaire was a contributor in 2004.

Main project activity

The principal focus so far has been a Grid-enabled compute and storage farm, the CERN opencluster. It is based on 100 of HP's dual processors Integrity servers, Intel's Itanium Processor Family (IPF) processors, Enterasys' 10-Gbps switches and a high-capacity storage system based on IBM's SAN FS storage subsystem. The 10 GB Ethernet technology is used for both the local-area and the wide-area connections. Oracle 10g is deployed as the relational database software.

Important results have been produced via participation in LCG data challenges and service challenges, as well as individual demonstrations, tests and benchmarks. Delivery of results has also taken other shapes, as in the porting and integration of the entire EGEE/LCG-2 middleware stack on the 64-bit Itanium platform. Furthermore, many requirement documents as well as functional assessment reports have been delivered to the partners. More information can be found in the CERN openlab annual report.
In most cases the work is carried out as a direct collaboration with the research and development arm of the partner concerned.

Human resources

Core manpower for openlab is provided by CERN, and includes a full-time openlab CTO, an openlab manager and administrative support, as well as contributions to specific tasks from a wide range of technical experts within CERN’s IT Department. In addition to this, the most important manpower contribution comes from the employment of post docs (CERN fellows) who are dedicated to the openlab programme of work. These positions are mainly financed by the partners and employed by CERN. During the last two years the CERN openlab has also had the privilege of employing Marie Curie fellows sponsored by the EU. These fellows have worked on the same challenges as the post docs under a regime that mixes skill deployment with comprehensive learning.

Quite naturally, the openlab provides an excellent technological framework for CERN technical students who wish to complete a master thesis in one of the established fields of technology. Additional human resources are found in a dedicated CERN openlab student programme, permitting a substantial integration of summer students. This programme has repeatedly employed a dozen or more summer students with a large fraction directly related to the opencluster projects. Some of the students have been co-funded by openlab partners.

Additional activities

In brief, CERN openlab has also acted as a demo centre, a forum for workshops on advanced technological topics, and a channel for sponsorship of events, such as the Computing in High Energy Physics (CHEP) Conferences.

Proposal for projects in "CERN openlab II"

Principal mission

The years 2006 – 2008 correspond to the period of the full deployment of the LHC Computing Grid. Real data will start flowing from the LHC detectors around the middle of 2007 but service challenges will simulate the real load on the Grid right from the start. The need to understand and master the continuous technological evolution will be just as important during this period as during the current three-year period. A natural option in this context is to build upon the expertise and reputation established by openlab I.

The mission of openlab II will continue to be a close collaboration with major trustworthy industrial partners able of demonstrating technological leadership and delivering viable solutions. In addition, closer collaboration with the European Grid infrastructure project EGEE (Enabling Grids for E-sciencE) will be formalized, by taking advantage of the fact that EGEE enters a second phase in May 2006, called EGEEII. EGEE is currently the largest multiscience Grid deployment project worldwide, with 70 institutional partners in Europe, the US and Russia.

Main project activities

Based on an openlab workshop held in June 2005 on "Industrialising the Grid", a number of topics were defined of interest to CERN, the EGEE project, and the openlab partners. The resulting shortlist is included at the end of this document. At the annual openlab sponsors meeting in July 2005, two specific themes were singled out as being of particular interest to some of the current openlab partners. These are:

- A grid interoperability and integration centre with close links to EGEEII (referred to as GIC)
- A platform competence centre with main focus on the PC-based computing hardware and the related software (referred to as PCC)

GIC

The grid interoperability project is proposed as a reinforcement of the activities within EGEEII, and will allow the partners to take part in the integration and certification of grid middleware. Partners will be expected to contribute hardware for a local test-bed as well as suitable grid resources remotely.
The project will focus on the following activities:

- Testing and certification: This activity will encompass tests of the EGEE II stacks on test-beds provided by the partners.
- Support, analysis, debugging and problem resolution: This activity will depend on the problems encountered on the contributed test-bed.
- Interoperability: This activity will review current levels of interoperability, also with industrial middleware stacks proposed by the partners, with an aim to strengthen them and improve standardization of interoperable grid middleware in general.
- Capture of requirements: This activity will be based on all the other activities already mentioned.

The Appendix provides more detailed information on the GIC.

PCC

The platform competence centre will address a range of important fields, such as application optimization and platform virtualization:

- Software and hardware optimization is seen as a vital part of the Grid deployment, since the demand for resources by the scientists is very likely to outstrip the available resources, even inside the Grid. Such optimization relies on deep knowledge of the architecture of the entire computing platform. On one hand this covers hardware items, such as processors, memory, buses, input/output channels, on the other hand it covers the ability to use proficiently advanced tools, such as profilers, compilers and linkers, specially optimized library functions, etc.
- Platform virtualization can be used to allow grid applications to enjoy a highly secure and standardized environment presented by a "virtual machine hypervisor", independent of all the hardware and software intricacies. The virtualization concept, initially offered as a pure software solution, will gradually move into hardware, allowing greatly improved performance of both the hypervisors and its "guests".

These two themes do not exclude other activities being launched in the openlab framework. Indeed, there are plans for two European SMEs to join as openlab contributors in an IT security project. The appendix provides a list of possible topics that other (probably of a smaller scale) projects would address.

Human resources

It is expected that most of the staff will be recruited as CERN fellows, but there will also be the option to employ individuals in LD (Limited Duration) positions. As in the past, openlab will offer the possibility for technical students or doctoral students to prepare their theses inside the projects. In a similar fashion, the summer student program will be continued as before.

A new feature of openlab will be the possibility for staff from the partners to spend a sabbatical period at CERN. The exact conditions for such sabbaticals will be decided on a case by case basis.
The CERN contribution of openlab core management and technical expertise will continue on the same level as for openlab I.

Additional activities

The existing activities under this heading will continue as before: CERN will act as a demo centre, a forum for workshops on advanced technological topics, and a channel for sponsorship of events. Public relations of CERN openlab has been emphasized by partners as an area they would like to invest in, and there are plans to integrate dedicated short-term communication interns at CERN, supported by partners. Other relevant activities will be integrated as needed.

Conclusion

During the first three years of existence, the CERN openlab has demonstrated the ability of CERN to work successfully with key industrial partners on the test and validation of new technologies. In the proposed second phase of CERN openlab, a formally defined relationship with EGEEII will increase the impact of the openlab partnership beyond LHC Computing to the wider domain of Grid-based scientific computing. Two projects for the second phase, a Grid interoperability centre and a platform competence centre, are at an advanced stage of planning with several existing partners, and there is scope for further projects.

Appendices

List of topics of interest for future collaboration

This list is constructed from the outcome of the CERN openlab workshop "Industrialising the Grid".

It provides topics for specific projects involving interested industrial partners. Such projects would be recognized as EGEEII associated projects as well as openlab projects. In the latter case, they would be viewed as smaller projects complementing the two major projects proposed for openlab-II.

1) Economic models of the Grid – what can EGEE operations teach us and how can industry help? Use EGEE to test future commercial solutions for offering Grid services
2) Job monitoring and costing and job performance, in particular metrics for comparing Quality of Service between Grid resources and conventional resources such as local clusters
3) Training and education in Grids, including web-based training (t-Infrastructure)
4) Commercial support of middleware to ensure long-term reliability by, for example, measuring and improving methods for automatic deployment of Grid middleware or providing a service for porting applications to the middleware
5) Simplified automatic management of resources on the Grid (autonomic computing technology)
6) Analyzing possible licensing models, and developing solutions that encourage the uptake of Grid technology by application vendors
7) Ensuring new releases of applications are compatible with middleware
8) Preparing for long-term sustainability of the infrastructure (analysis of how the private sector manages similar infrastructures, and resulting knowledge sharing)
9) Operating an industrial testbed for industrial applications, based on gLite (possible EU project?)
10) Test and validation of alternative (commercial, open source) solutions to current middleware (e.g.TomCat)

First Elements for a GRID Interoperability Centre

Overall objective

Create a multi-platform / multi-vendor testbed to test **interoperability** and **scalability** of Grid systems and evaluate Grid applications, with a particular focus on standardization.

Model

The test bed is conceptually formed of three layered components:

1. The **Grid Resource Centres** provided and operated by industrial partners and located of their premises as well as on the CERN site
2. The **Grid Intermediate layer**, formed of middleware components and other necessary components - such as monitoring and management modules
3. The **Grid Applications** which run on top of the middleware infrastructure.

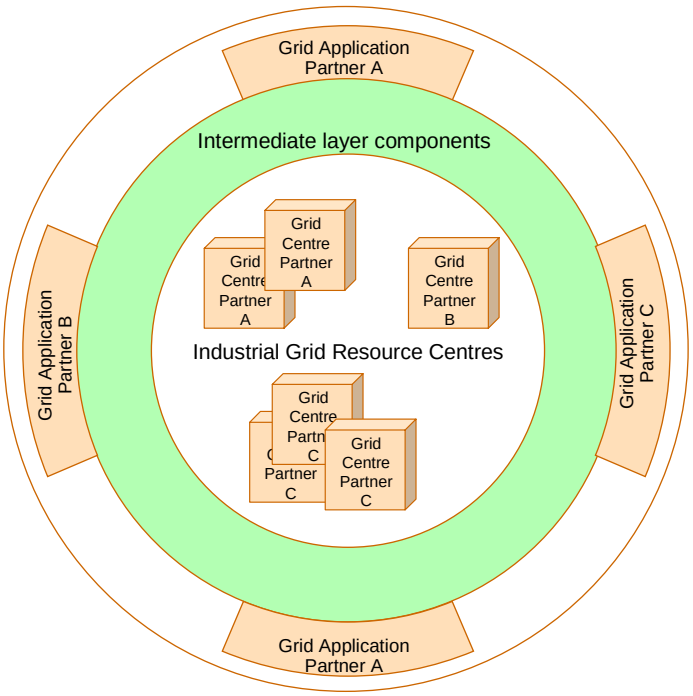

Figure 1 The three-layer model of the Grid Interoperability Centre testbed

The Grid Resources Centres

The Grid Resources provided by an industrial partner are not located on a single site. They are located:

- On one or several sites of the partner
- On the CERN site.

It is recommended that the Resource Centres located on partners' premises are formed of a collection of identical replicable modules. Each module would provide the full Grid functionality. Such a structure will facilitate the scalability of the testbed.

In such a case, each partner would provide on the CERN site at least one full-functionality module.

If partners only host resource modules of their own on their premises, this may be called the ***Combined-Resource model***, as exemplified in the figure below.

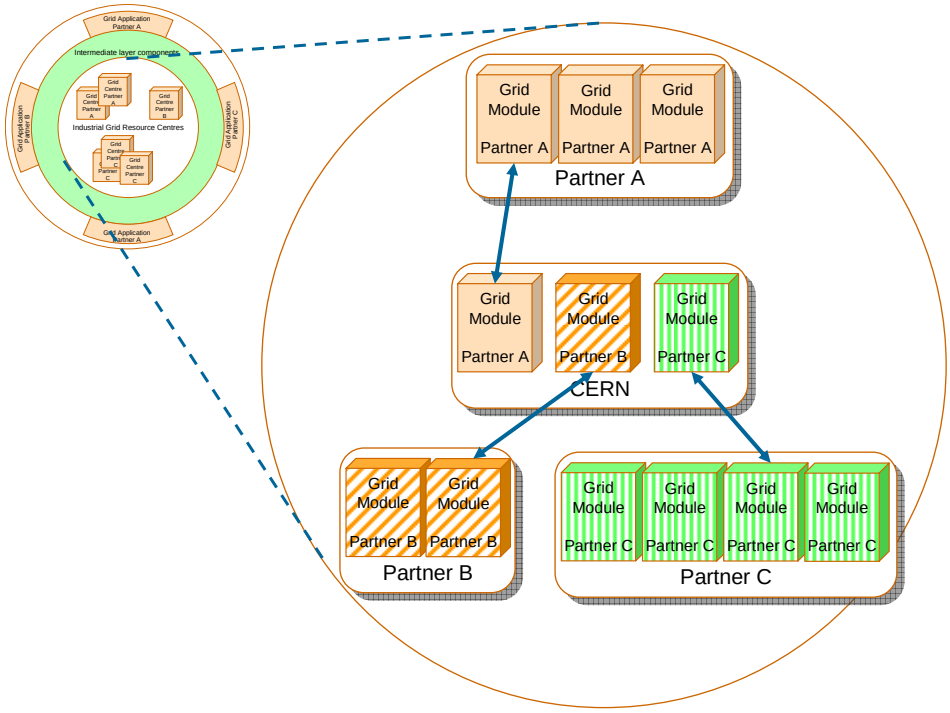

Figure 2 Logical Structure of the GIC Resource Centres – Combined-Resource Model

If, on the contrary, partners host modules from differing sources on their premises, this may be called the ***Shared Resource model***.

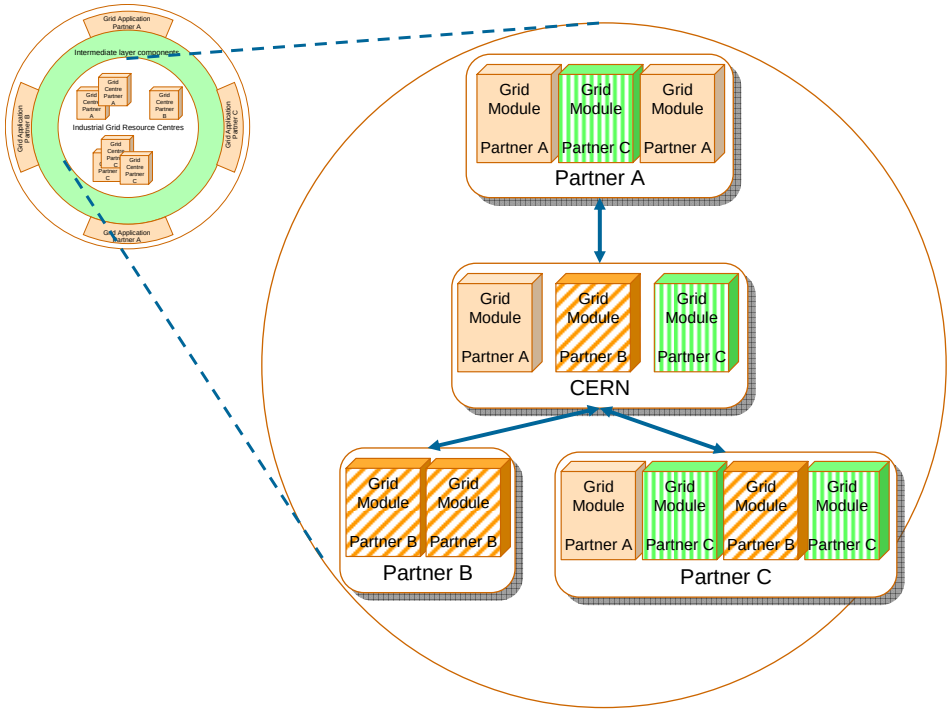

Figure 3 Structure of the GIC Resource Centres – Shared-Resource Model

Participants and roles

Role of Industrial partners

Industrial partners act at two levels:

1. Grid Resource Centres
    - They provide the Resource Centres
    - They operate the Grid Resource Centres on their premises, and provide resources for their operation on the CERN site at CERN.
2. Applications
    - They organize applications on top of the Intermediate Layer.

In addition, they provide resources for the operation of the Intermediate Layer components (middleware and other operational components). This may take the form of placing or providing funds for personnel on the CERN site.

Role of CERN

In addition to providing the overall project coordination, CERN makes available to the project the gLite middleware and other CERN Grid management tools.

CERN may also act as an application provider by running parts of the LCG particle physics applications. This may be useful to assess the option of the absorption of peaks of load by industrial Grids.

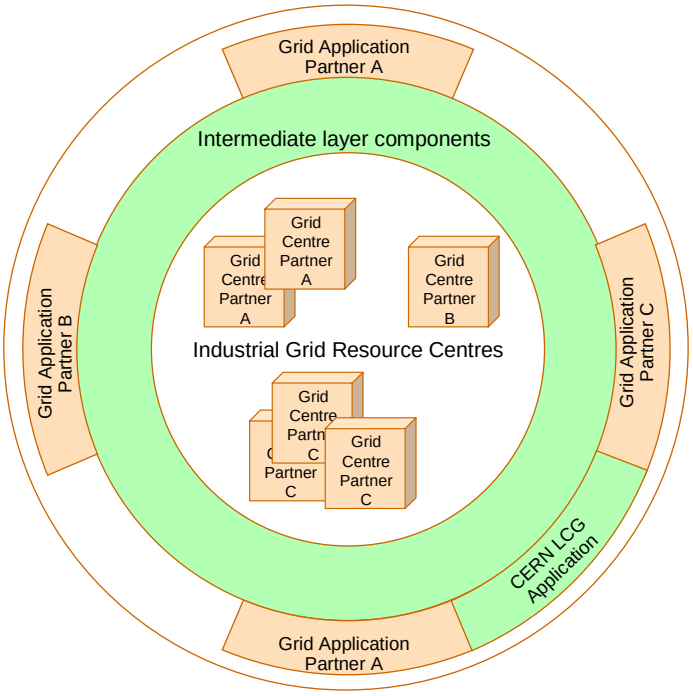

Figure 4: CERN as one provider of applications

Use of the industrial resources

The industrial partners may wish to control which resources are used by which applications. The VO principles as well as the authorization mechanisms provided by the middleware should allow maximum flexibility to organize the allocation and sharing of resources.

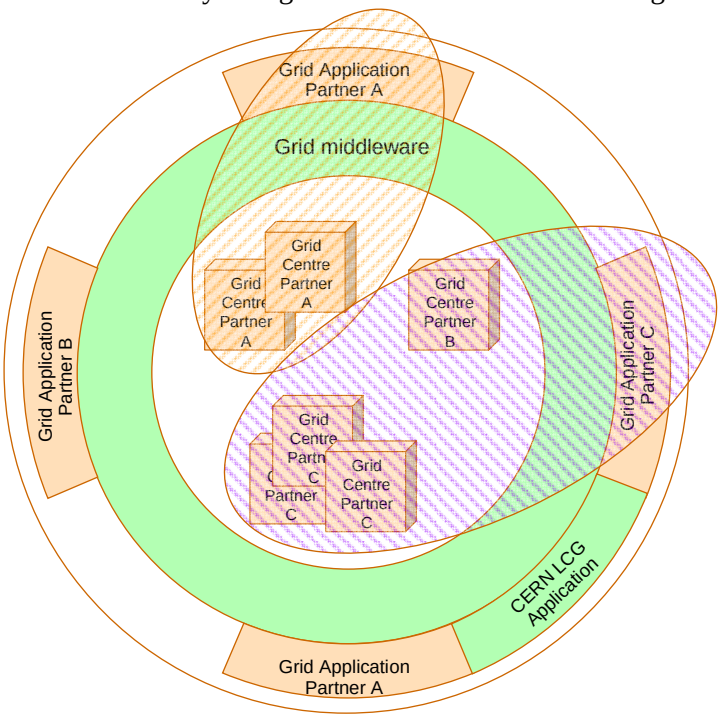

Figure 5: VOs to control access of applications to Resources Centres

Specific objective

Initial list:

- Validate **Interfaces** between
  - Industrial Grid Resources Centres and middleware
  - Middleware (and other operational components) and industrial applications
- Test **Interoperability** between Grid industrial implementations of Grid **open** standards.
- Create expertise on **industry-class Grid operation** of Grid infrastructures
- Study Grid services **cost models** and provide a testbed for Grid services business models (in particular through CERN applications)

Relationships with EGEE II

The project is referred to in the EGEE II proposal as one dimension of the industrial dissemination plan.