

February 2005

Grid Computing, E-Science and Applications in Industry

Urs Hoffmann
CERN

Follow this and additional works at: <http://aisel.aisnet.org/wi2005>

Recommended Citation

Hoffmann, Urs, "Grid Computing, E-Science and Applications in Industry" (2005). *Wirtschaftsinformatik Proceedings 2005*. 90.
<http://aisel.aisnet.org/wi2005/90>

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik Proceedings 2005 by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

In: Ferstl, Otto K, u.a. (Hg) 2005. *Wirtschaftsinformatik 2005: eEconomy, eGovernment, eSociety*;
7. Internationale Tagung Wirtschaftsinformatik 2005. Heidelberg: Physica-Verlag

ISBN: 3-7908-1574-8

© Physica-Verlag Heidelberg 2005

Grid Computing, E-Science and Applications in Industry¹

Hans Hoffmann

CERN

1 Introduction

The early universe was extremely dense, hot and opaque for electromagnetic radiation until it cooled off and expanded sufficiently to become transparent to light and astronomical observation. Particle Physics creates matter at temperatures or energies that correspond to conditions in the universe at minute fractions of a second after the big bang and not observable by any other means. In turn such very high energies correspond to extreme short wavelengths of the matter particles involved and allow exploring at very small scale the quarks and leptons, the constituents of today's periodic table of elements, named the "Standard Model" and the forces that describe the interactions between them. Questions asked today are the origin of mass, how to reconcile quantum mechanics describing the microcosm and relativity describing the macrocosm (keyword "string theories") and a unified theory describing the variety of particles and forces.

CERN was founded 50 years ago to provide the European technical and technological infrastructure for the collaboration of its Member States towards exploring such fundamental questions.²

Particle Physics has extended its reach into the microcosm by an order of magnitude every 10 to 15 years over the past 40 years corresponding in turn to order of magnitude higher energies for reactions between elementary particles. Important steps were the discovery of the antiproton³, the J/Ψ ⁴ and the W, Z particles⁵ (No-

¹ Numerous colleagues have contributed to the content of these notes, in particular persons from the European DataGrid (<http://cern.ch/edg>), the LHC Computing Grid, LCG (<http://cern.ch/leg>), and the EGEE (<http://cern.ch/egee>) projects. At the respective web sites ample and comprehensive information can be found. I would like to mention in particular Fabrizio Gagliardi, CERN, Les Robertson, CERN, Manuel Delfino, now at IFAE-PIC, Barcelona, Alois Putzer, Kirchhoff Institut, Heidelberg and A Reinefeld, Zuse Institut, Berlin. All of those have significantly contributed to my understanding of the subject and, more importantly, to the mentioned projects.

² General information on CERN can be found in the CERN web-pages under <http://www.cern.ch>

³ <http://www.nobel.se/physics/laureates/1959/>

⁴ http://www.bnl.gov/bnlweb/history/Nobel/Nobel_76.html

bel Price for Physics 1984 for the CERN physicist Carlo Rubbia and the CERN engineer Simon van der Meer) and the precise determination of the standard model parameters. CERN's LHC project⁶, to become operational in 2007, will increase this range with respect to the Tevatron at FNAL near Chicago⁷ again by an order of magnitude to around 10-18m and to objects of up to 2000 GeV mass. It will exceed the discovery potential of the original CERN Proton Synchrotron accelerator, firstly operate in 1959, measured in cms energy, by a factor of 2000.

The statistical nature of high energy particle collisions requires extreme selectivity to identify the rare events of interest which occur at levels of 1:10¹² collisions or even less frequently for interesting physics at the CERN LHC. This denotes the other challenge of particle physics, namely to identify ever more rare events with ever more sophisticated detection and identification means. An eminent recognition of CERN's contribution to detectors is the Nobel Price for Physics for the CERN physicist Georges Charpak for his development of multiwire proportional chambers in 1992⁸. The most prominent role amongst the technologies used are, however, Information and Communication Technologies, ICT. Also there CERN has made significant contributions as recognised by the Millennium Price of Finland given last year to Tim Berners Lee for his development of the WWW⁹ at CERN about 15 years ago.

Whereas the early discoveries, order of 50 years ago, were still performed with, for example, photographic observation and recording of individual events now electronic detection, recording and analysis of large numbers of events is customary. Increasing amounts of detailed parameters are measured and online selection and filtering is used, followed by sophisticated offline reconstruction, simulation and analysis programs. The LHC general purpose experiments today, ATLAS and CMS, have in excess of 100 million electronic channels and are able to digest $\sim 10^9$ events/s, each with an average multiplicity of ~ 100 generated secondary particles.

In fact, the influx of data into ATLAS¹⁰ or CMS at preamplifier level and assuming everything was digitized runs at an equivalent PetaByte/s in-flow of data of which ~ 100 MByte/s are recorded offline, or ~ 100 events/s. This corresponds to an online selection and reduction of the available data by a factor of 10^7 . The recording rate for the LHC heavy ion detector reaches 1.25 GBytes/s originating from the extremely high multiplicity of the individual ion-ion collisions.

⁵ <http://www.nobel.se/physics/laureates/1984/rubbia-lecture.pdf>

⁶ <http://lhc-new-homepage.web.cern.ch/lhc-new-homepage/>

⁷ <http://www.fnal.gov/>

⁸ <http://nobelprize.org/physics/laureates/1992/>

⁹ http://www.technologyawards.org/index.php?technologyawards=1104743195b0c38a534b13af7d86125023aef161cc&article_id=3338&_from_id_=2912

¹⁰ For all LHC experiments compare: <http://greybook.cern.ch/>

CERN and its collaboration institutes perform a continuous IC Technologies watch named “PASTA”¹¹ in the fields of semiconductor technology, secondary storage, mass storage, networking-, data management - and storage management technologies and high performance computing solutions. The important motivation for these studies is to identify the best cost/performance solutions.

A final important ingredient to the contribution from a fundamental science to discussions on technology and innovation comes from the collaborative nature of particle physics. At this time CERN collaborates with almost 300 institutes and more than 4000 scientists in Europe and in excess of 200 institutes and another almost 2000 scientists elsewhere in the world. The ATLAS or CMS collaborations¹² are about one third of this each and amongst the largest long-term, global scientific collaborations ever existing until today, with well defined and commonly pursued scientific objectives, each realising devices worth in excess of half a billion € each, CERN contributing ~20%. In the grid context we shall name such collaborations “virtual organisations”, distant groups of people collaborating coherently towards a common objective.

2 Grid Computing to Solve the LHC Worldwide Data Analysis Challenge

The Technical Proposals of the LHC experiments, the first comprehensive papers describing the complete scope of the detectors and published in 1994¹³ gave little information on computing and application software issues. It seemed possible to solve the technical problems but the development of the corresponding technologies was still to far off the level required. (Moore’s law)

Following the approval of the LHC project, accelerator and the four experiments (ALICE, ATLAS, CMS, LHCb) in 1996 computing needs started to be considered again. In a common effort between ATLAS and CMS the MONARC project¹⁴ provided key information on the design and operation of the worldwide distributed Computing Models for the LHC experiments.

These computing models envisaged thousands of physicists involved in data analysis at institutes around the world making use of a complex, tiered set of compute and data servers connected by powerful wide-, regional- and local area networks.

¹¹ <http://lcg.web.cern.ch/LCG/PEB/PASTAIII/pasta2002Report.htm>

¹² <http://cern.ch/atlas>; <http://cern.ch/cms>

¹³ CERN/LHC/94-43; LHCC/P2, 15 December 1994, compare CERN document server

¹⁴ <http://www.cern.ch/MONARC/> and “Models of Networked Analysis at Regional Centres for LHC Experiments” (MONARC), phase 2 report; KEK preprint 2000-8, CERN/LCB 2000-001, April 2000, M Aderholz(MPI) et al.

Each of the experiments envisaged storing and partially distributing data volumes of Petabytes/y. The MONARC project provided for a simulation tool to assess and optimise the cumulated performance of the distributed facility.

In the summer and fall of 1999, it was then realised that Computational Grid technology¹⁵, conceived for coupling supercomputers to achieve maximum performance supercomputing, could be extended and adapted to data-intensive, high throughput tasks serving a worldwide community.

It was also realised that important amounts of R&D were required to develop the workflow and resource management tools needed to manage such a worldwide, distributed, multi-user and high throughput “Data Grid” system. Two activities were launched at that time, the LHC Computing Review and the formulation of “data grid” R&D proposals with EU and US funding agencies.

2.1 The LHC Computing Review¹⁶

The Computing Review evaluated the current situation, plans and prospects of data management and computing at the LHC. The Review involved representative software and computing experts from the experiments and the Information Technologies Division. The Review was supposed to establish:

- Solid requirements figures and identified common efforts between the partners.
- A proposal of how to organise the efforts.
- The need to collaborate closely, the experiments amongst themselves, with the Information Technologies Division at CERN and with national computer centres.

The Executive Summary of the Review stated: “The requirements to ensure the storage, management, simulation, reconstruction, distribution and analysis of the data of the four LHC experiments (ALICE, ATLAS, CMS and LHCb) constitute an unprecedented challenge to the High Energy Physics (HEP) and Information Technology (IT) communities. After critical assessment, the review accepted the scale of the resource requirements as evaluated together with the four experiments, namely a total of about 8000 k SI-95 as total processor power required, corre-

¹⁵ I. Foster and C. Kesselmann, *The GRID: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, San Francisco, 1998. Compare also: UNICORE (<http://www.fz-juelich.de/unicore/>) and e-Science Programme in the UK (<http://www.escience-grid.org.uk/>)

¹⁶ CERN/LHCC/2001-004, CERN/RRB-D 2001-3; 22 February 2001, “Report of the Steering Group of the LHC Computing Review” (S Bethke, (chair) MPI Munich, M Calvetti, INFN Florence, H F Hoffmann, CERN, D Jacobs, CERN, M Kasemann, FNAL, D Linglin, IN2P3)

sponding at the time (year 2000) to the cumulated compute power of about 300 000 "then" PCs, distributed into a world-wide system, the main parts to be interconnected by WANs with bandwidths in the 10 Gbit/s range and with a total data storage capacity on tape of ~30 Petabyte/year of LHC operation and on disk of ~10 Petabyte.

A multi-Tier hierarchical model should be the key element of the LHC computing model. In this model, for each experiment, raw data storage and reconstruction will be carried out at a Tier0 centre. Analysis, data storage, some reconstruction, Monte-Carlo data generation and data distribution will mainly be the task of several (national or supra-national) "Regional" Tier1 centres, followed by a number of (national or infra-national) Tier2 centres to which the end users would be connected appropriately. The CERN-based Tier0+Tier1 hardware for all LHC experiments should be installed as a single partitionable facility.

Grid Technology will be used to attempt to contribute solutions to this model that provide a combination of efficient resource utilisation and rapid turnaround time."

The Review requested further activities. In more detail the following points were raised:

- Data Challenges of increasing size and complexity should be performed as planned by all the experiments until LHC start-up.
- CERN should sponsor a coherent programme to ease the transition of the bulk of the physics community from Fortran to Object Oriented (OO, C++) programming.
- Further identified areas of concern were the limited maturity of current planning and resource estimates, the development and support of simulation packages and the support and future evolution of analysis tools.

Concerning resources the complete distributed facility was estimated in terms of materials cost and of personnel and material cost at CERN using the PASTA extrapolation. The CERN material effort was about one third of the total distributed facility.

The approach to Maintenance and Operation of the LHC computing system included the strategy of rolling replacement within a constant budget.

The construction of a common prototype of the distributed computing system should be launched urgently as a joint project of the four experiments and CERN/Information Technologies department, along with the major Regional Tier Centres. It should grow progressively in complexity and scale to reach ~50% of the overall computing and data handling structure of one LHC experiment in time to influence the acquisitions of the full-scale systems, corresponding to about 20% of the total.

2.2 “Data Grid” Projects

The vision of Grids is the long standing wish to connect computing resources together in a transparent way when they are required for a given task, or expressed in the words of Larry Smarr and Charles Catlett, 1992: „Eventually, users will be unaware they are using any computer but the one on their desk, because it will have the capabilities to reach out across the internet and obtain whatever computational resources are necessary”

In the layered view of wide area networking services originally the following layers could be associated with an integrated science environment (snapshot from DataGrid):

- “Problem Solving Environment” layer: domain specific application interfaces for scientists
- Grid Middleware layer:
 - Grid Application: job-, data- and metadata-management
 - Collective Services: grid scheduler, replica manager, information and monitoring services
 - Underlying Grid Services: SQL database-, computing element-, storage element services, replica catalogue, authorisation, authentication, accounting, service index
 - Fabric services: resource-, configuration-, node and installation- and storage managements as well as monitoring and fault tolerance
- Networking and Computing Fabric layers

The challenge of constructing the middleware is to identify an architecture of the tasks which allows factorizing them into individual services whilst preserving simplicity, overall efficiency and transparent use. This is the reason for CERN and its collaborating institutes to foresee ample testing opportunities involving the whole community.

Several Data Grid projects were launched subsequently covering different aspects of the problem. In Europe and supported by the EU DataGrid (large datasets, high throughput, many users), DataTag (quality of service and high bandwidth in grids), crossgrid (more applications) and EGEE (grid infrastructure to enable e-science in Europe); in European nations, to name some of the major ones, in which CERN was involved directly.

The European DataGrid¹⁷ project was elaborated, proposed and approved in the years 1999 and 2000. It was funded by the European Union and the participating

¹⁷ <http://cern.ch/edg> and using a short description from EDG

institutes with the aim of setting up a computational and data-intensive grid of resources for the analysis of data from scientific exploration.

The main goal of the DataGrid initiative was to develop and test the technological infrastructure that will

- enable the implementation of scientific “collaboratories” or “virtual organizations” where researchers and scientists will perform their activities regardless of geographical location and where they will collaborate and share resources and facilities for agreed common objectives.
- allow interaction with colleagues from sites all over the world as well as the sharing of data and instruments on a scale not yet previously attempted.
- devise and develop scalable software solutions and testbeds in order to handle many PetaBytes of distributed data, tens of thousand of computing resources (processors, disks, etc.), and thousands of simultaneous users from multiple research institutions.

The DataGrid initiative was led by CERN together with five other main partners and fifteen associated partners. The project brought together the European Space Agency (ESA), France's Centre National de la Recherche Scientifique (CNRS), Italy's Istituto Nazionale di Fisica Nucleare (INFN), the Dutch National Institute for Nuclear Physics and High Energy Physics (NIKHEF) and UK's Particle Physics and Astronomy Research Council (PPARC) plus further fifteen associated partners come from the Czech Republic, Finland, France, Germany, Hungary, Italy, the Netherlands, Spain, Sweden and the United Kingdom.

The project covered three years, from 2001 to 2003 plus 3 month to wrap-up and consolidate the results, with over 200 scientists and researchers involved. The EU funding amounted to 9.8 million euros, mostly for people. The required matching contributions from the partners exceeded that amount.

The first and main challenge facing the project was the sharing of huge amounts of distributed data over the currently available network infrastructure. The DataGrid project relied upon emerging Computational Grid¹⁸ technologies.

The DataGrid project was divided into twelve Work Packages distributed over four Working Groups: Testbed and Infrastructure, Applications, Computational & DataGrid Middleware, Management and Dissemination. Figure 1 below illustrates the structure of the project and the interactions between the work packages.

¹⁸ <http://www.globus.org>

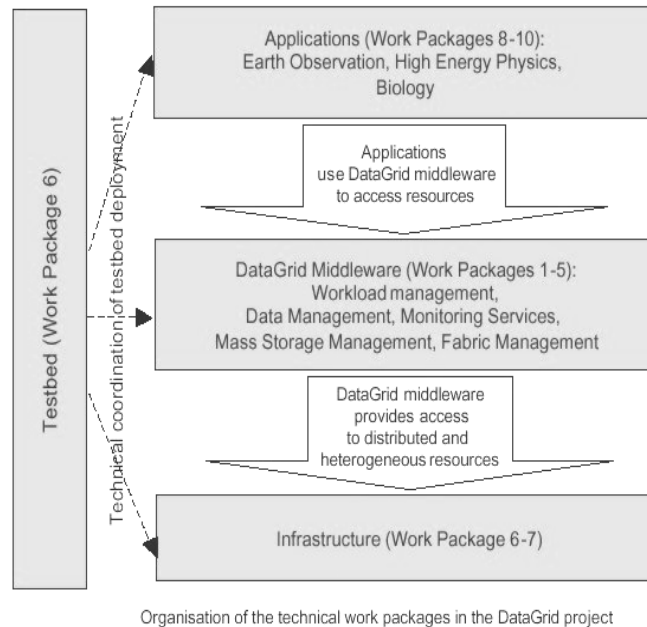


Figure 1: Interaction between Work Packages

More details on the functions and relations of the various middleware packages can be found in the DataGrid tutorials¹⁹.

DataGrid has obtained most of its objectives by the end of 2003 and terminated end of March 2004. The important “deliverables” were:

- Proof of viability of the grid concept,
- Provision of training (>200 persons trained) for the applications addressed, namely for particle physics, bio-medical and earth observation communities, based on extensive test beds,
- Delivery of a tested and stable middleware package for the LHC Computing Grid, the delivered package containing basic elements of the Globus toolkit, contributed by Trillium issuing the “virtual data toolkit” (VDT) together with Globus and a substantial number of components delivered by DataGrid.

In 2003 and within the now active 6th Framework Programme of the EU the possibility arose of a new grid project named EGEE²⁰ standing for “Enabling Grids for E-Science in Europe”. A number of the DataGrid participants and many further institutes from all over Europe joined in elaborating the proposal. With its ap-

¹⁹ <http://eu-datagrid.web.cern.ch/eu-datagrid/Tutorial/tutorial.htm>

²⁰ <http://cern.ch/egee>

proval for initial two years by the EU in 2003 the project has started in April 2004 and is the currently operating Grid infrastructure project in Europe.

It is complemented by other grid projects (DEISA, SEEGrid, DILIGENT and GN2)²¹ and supposed to provide some infrastructure for them. The most important of those is DEISA²², a grid of major European supercomputer centres.

The desired goal of the EGEE project is to create a general European grid infrastructure of production quality on top of present and future EU Research and education networks. The project builds on the EU and EU member states major investment in Grid Technology and exploit the international connections the US and Asia Pacific. It is, in particular, accepted to use the particle physics community as the early adopters and for most of the initial tests and debugging to be performed during the first two years of the project. To this end, the project foresees a close collaboration with the LHC Computing Grid Project described below.

The basic middleware functionality can be summarized as job management services (submit, monitor, control), resource management services (location, matching), data management (replication, location), security (services as they evolve) and operations and fault diagnosis (accounting, logging). An overview of that basic functionality of the grid middleware is given in Figure 2 below

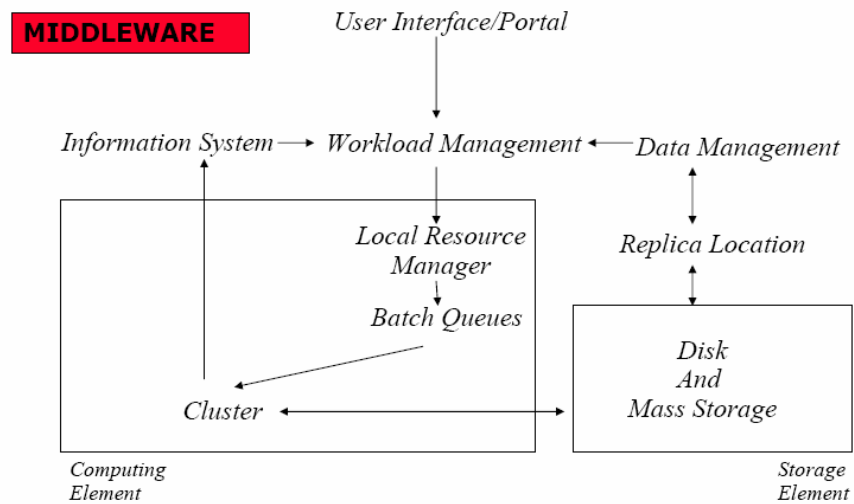


Figure 2: Functional overview of a (Data-) Grid

²¹ <http://public.eu-egge.org/conferences/2nd/> providing a very complete set of information on the current status of the project

²² <http://cern.ch/datatag>

2.3 Networking Efforts

In the past 20 years networks for digital information transfer and exchange have had a phenomenal increase of available bandwidth from then customary ~10Kb/s to today's equally customary 10 GB/s, at least between privileged nodes, an increase by a factor of 106. More importantly during that period digital networks have become a generally accepted commodity. The European GÉANT connections between the National Research and Education Networks (NRENs) in Europe is probably the world's leading research network, connecting more than 3100 universities and R&D centres in over 32 countries across Europe and providing connectivity to most other parts of the world.

However, grids pose particular requirements to "quality of service" of networks and the EU and US funded DataTag²³ project addressed those.

The DataTAG project was to create a large-scale intercontinental Grid testbed that focuses upon advanced networking issues and interoperability between these intercontinental Grid domains, hence extending the capabilities of each and enhancing the worldwide programme of Grid development.

The project addresses the issues which arise in the sector of high performance grid and inter-grid networking, including sustained and reliable high performance data replication, end-to-end advanced network services, and novel monitoring techniques. The project has set the "internet land speed record" for data transmission between CALTECH and CERN several times.

The DataTag project is now terminated and its tasks promoted further by EGEE and GEANT.

2.4 The LHC Computing Grid²⁴ Project (LCG)

In September 2001 the CERN Council approved the first phase of the LHC Computing Grid project based on the needs established in the Computing Review discussed above and aimed at full exploitation of LHC data by all participating institutes and physicists. The Council followed the proposal of the CERN Management of a two-phased LHC Computing Grid Project resulting from discussions with the experiments and collaborating institutes subsequent to the review, and its proposals for action. It was stated in more detail that

- the size and complexity of this task required a worldwide development effort and the deployment of a worldwide distributed computing system, based on

²³ <http://www.deisa.org/>

²⁴ <http://cern.ch/lcg>

tiered regional centres, to which all institutes and funding agencies participating in the LHC experimental programme were asked to contribute.

- the envisaged common exploitation by the LHC experiments of the distributed computing system required coherence, common tests and developments, and close co-operation between all nodes of the computing system amongst each other and in particular with the CERN installations.

The LHC Computing Grid Project was conceived as a two-phased approach to the problem, covering the years 2001 to 2008 or in more detail:

- Phase 1: Development and prototyping at CERN and in Member States and Non Member States from 2001 to 2006, requiring expert manpower and some investment to establish a distributed production prototype at CERN and elsewhere that would be operated as a platform for the data challenges of the experiments.
- Phase 2: Installation and operation of the full world-wide initial production Grid system in the years 2006 to 2008, requiring continued manpower efforts and substantial material resources.

Further a formal project structure was proposed that would ensure the achievement of the required functionality and performance of the overall system with an efficient use of the allocated resources. Appropriate liaisons with existing grid projects as well as with industry would be put in place to promote the efficient use of resources and avoid duplication of work.

Given the special roles as host lab, as the Tier-0 and Tier 1 centre where storage of the raw data emerging from the four detectors and their first pass reconstruction would be performed it was proposed that CERN should lead the project and to perform the necessary CERN-based development and prototyping activities.

The LCG project aims at creating a global virtual computing centre for particle physics, focusing at the service aspects for LHC physics and its world-wide community of physicists.

LCG has four working areas, namely the

- application area for the LHC experiments, including the development of the software environment, the joint projects between the experiments, the data management and the distributed analysis,
- grid deployment area, establishing and managing the Grid Service, middleware certification, security, operations, registration, authorization and accounting, to be performed together with the EGEE project described above and possibly other complementary grid projects elsewhere in the world,
- grid technology or middleware area, providing a base set of grid middleware – acquisition, development, integration, testing and support, acquiring software from grid projects or industry,

- fabric area with large compute cluster management, data recording, cluster technology, networking and computing service at CERN.

The latter three areas would proceed gradually offering increasingly demanding services, starting with simulation, then batch analysis and later interactive analysis. Middleware acquired from DataGrid would be used first and then re-engineered middleware from projects such as EGEE or others based on increased experience with test beds.

LCG 1, the first LCG production grid service, 24hrs/day and 7 days/week, has become operational in March 2004 based on DataGrid, US-Grid 3 and Nordugrid middleware and was tested with "data challenges" or basically batch jobs by the experiments experiments. Releases with upgraded functionality succeeded each other until summer 2004 where LCG 2 offered a more stable environment for the planned more ambitious further data challenges of the LHC experiments.

In the 4th quarter of 2004 the project has seen the continued growth of the LCG-2 infrastructure (which also forms the initial EGEE production service), to more than 80 sites, providing more than 9000 CPU. The charts show the build up of resources during the quarter, and the numbers of jobs running in the service. The peak of jobs was during the final stages of the data challenges in August.




Figure 3.a: Total and free numbers of CPU, operated under LCG 2(EGEE) from July to September 2004




Figure 3.b: Running and waiting jobs, operated under LCG 2(EGEE) from July to September 2004

The overall efficiency of jobs was of the order of 65-70% including all sources of inefficiency and failure. A significant amount of jobs failed owing to lack of preparedness by some of the 80 computer centres participating.

For the Grid Technology area the important event in autumn 2003 was the publishing of the ARDA (Architectural Roadmap Towards Distributed Analysis) study. The results of this study were to give rise to new prototype grid middleware testbeds, based on "grid services", a variant of the much advertised web-services.

The prototype middleware allowing for a distributed, "chaotic" analysis has been named gLite for "lightweight middleware for grid computing". The requirements are to be quickly and easily deployable, to use existing services where possible for re-engineering, to interoperate with other middleware packages, resilience and fault tolerance, co-existence with deployed infrastructure i.e. to run as application on LCG 2 and the US-Grid 3. The prototype should become operational now, early 2005.

LCG and EGEE are at present engaged in testing usable grid middleware components approaching production level. In parallel major efforts are ongoing developing the European and world-wide computing capacities and the readiness of computer centres to operate in grid mode. The particle physicists are engaged in finishing off their very large software packages and to make them easily operational for the non-specialist user interested in data analysis in a grid environment. Therefore here basically only a status report is given concerning first operational experience in using grid-like distributed computing services. A corresponding talk in a year from now would be able to describe first really operational middleware packages and demonstrate first solutions to Grid Security problems that will be addressed in 2005.

3 Particle Physics and E-Science, E-Infrastructure, Cyberinfrastructure

Science that is competitive at a global scale will require coordinated resource sharing, collaborative processing and analysis of huge amounts of data produced and stored by many scientific laboratories belonging to several countries and institutions. Such sciences have complete mechanisms of collecting the relevant raw data, turning raw data into validated, calibrated pieces of information. Such pieces of information are ordered, completed and validated and described coherently. Finally in the comparison with underlying theoretical understanding such data confirm or not the current understanding and complement and extend the knowledge produced by that science. With the exponential growth of the amounts of data in many sciences this process is performed relying on ICT.

Sciences which analyse and understand all their data in this manner are "e-sciences". "E" is standing for enhanced, empowered, enabled or brought to full potential.

In the words of the former Director of the UK Research Councils, Sir John Taylor, “e-science is about global collaboration in key areas of science and the next generation of infrastructure that will enable it”²⁵.

Particle physics in the LHC context is probably the most advanced e-science at this time or an “e-science par excellence”.

In many countries grids and e-science are seen as the next step towards a comprehensive “cyber infrastructure” as the environment is named by the US-National Science Foundation. They state that the cost and complexity of the 21st century science requires the creation of an advanced and coherent global information infrastructure (info structure). “Info structures” will be driven by global applications changing the ways science is done. Particle physics with the LHC challenge is such an early driving science and the US is joining forces to globally advance the “info structures”.

In a similar way and possibly even in advance of the US the EU has made e-infrastructure one of their key actions in the 6th Framework Programme, supporting early production grids such as EGEE as e-infrastructure and the “complex problem solving” initiative to advance the scope of the comparatively simple particle physics to more complex sciences and their enabling application technologies and underlying advanced architecture, design and development of the next generation grids.

The UK e-science initiative, however, was the instrumental to advance massively application collaborations in the e-science sense and serves as model for Europe and its individual nations. CERN and the LHC community are grateful to their substantial support as well as to the German BMBF, the Italian INFN, the Spanish support and to many other nations or institutions helping with the progress.

4 Interest of E-Science and Grid Computing for Industry

The path to involve industry in grid computing can first pass through sciences enabled by full use of ICT and grid infrastructures and providing novel business opportunities.

The other path is that industries see advantages in dynamically linking activities across the particular firms in order to propose novel competitive services or devices beyond the reach of the individual firm.

²⁵ E-Science Programme in the UK (<http://www.escience-grid.org.uk/>)

In science the path will invariably involve collaboration of heterogeneous and distant institutes combining their efforts to address problems that they could not solve individually and where funding agencies would not and should not be willing to complement their resources to enable them individually.

Collecting ad hoc critical masses to address novel problems is the name of the game.

The collaborative nature of world-wide particle physics is a good example for using limited, however world-wide resources to address the most relevant questions of the science. Often the limited resources are the competent people to address all problems posed. The pervasive technology required to make up e-science efforts are ICT.

In Germany the D-Grid initiative²⁶ expresses a "Vision of a new form of scientific activity: e-Science", stating that

- the modern Knowledge and Information Society uses the multiple possibilities of efficient communications, easy access to large data and information repositories and powerful computers.
- these new methods offer the opportunity to obtain qualitatively and quantitatively better results in science and industry. However, full use of the offered opportunities requires mastering the distributed, dynamically linked system components.
- **virtualisation** is therefore the key notion to future Grid systems. In contrast to existing distributed information systems Grid-Systems of the next generation virtualise data-, information- and computation-services by hiding technical details of the concrete distributed realization behind interfaces. The user only sees the "socket in the wall" giving access to the wanted services and should not be aware of the technical details necessary to obtain computation-, data- or application services. Future grid-systems will enclose thousands of geographically distributed, heterogeneous resources linked together by wide area networks such as the internet. They will be autonomous, self-supervising and self-configuring to overcome errors and lack of resources. Further they will be secure, reliable, resilient, universally accessible and offer high level services. In addition they will be completely transparent with assured authentication and authorization. The user need not care about the resources or about their location. Like embedded microcomputers in household devices, cars or mobile communication systems the resources of future grid systems will be transparent and only the provided services appear.
- apart from the virtualization the **cooperative or collaborative use** of computer resources will be a key element of future grid generations. This particular fea-

²⁶ <http://www.d-grid.de/>

ture will enable completely new forms of collaborative work in science and industry. Based on particular scientific, but geographically or even internationally distributed facilities grids will support the creation of virtual organizations of "critical mass" enabling the collaborating participants to use the facilities to solve novel and common scientific problems that could not be tackled individually.

- this will allow in particular interdisciplinary, European or global collaborations at a new and advanced level.
- distributed industrial product development will be enabled as well as spontaneous clustering of scientists for a particular scientific purpose such as a common data analysis by distant but most qualified scientists as well as modeling of novel structures, visualization or common publication.

This means a change of paradigm in value creation in industry or science.

The European DataGrid and EGEE projects are both meant to be multi-purpose projects based on a specifically interested and driving community, here particle physics and CERN. They both invite other sciences and, in particular, industries to participate with (test-) projects of their own to make use of the novel infrastructure. To this end they both had or have an Industry Forum with the purpose:

- The role of the Industry Forum in the Enabling Grids for E-Science in Europe (EGEE) project is to raise awareness of the project amongst industry and to encourage businesses to participate in the project. This will be achieved by making direct contact with industry, liaising between the project partners and industry in order to ensure businesses get what they need in terms of standards and functionality and ensuring that the project benefits from the practical experience of businesses.
- The members of the EGEE Industry Forum are companies of all sizes who have their core business or a part of their business within Europe.
- The Industry Forum will be managed by a steering group consisting of EGEE project partners and representatives from business.

In the EGEE Industry Forum²⁷ the intention is therefore to listen to requirements of industries and to promote uses of the infrastructure. The means will be to provide consultancy, training and collaboration when appropriate. This is in particular true for the technology minded industries of all the small and medium size industries in Europe.

The UK e-Science programme provides for numerous examples for such activities and collaborates directly with EGEE and the EU.

²⁷ <http://public.eu-egee.org/industry/>

The D-Grid initiative foresees some funding for comparable initiatives and projects.

The most important ingredient of an initiative is the clearly perceived added value from collaborating with an external partner towards a novel purpose, service or product.

5 Fundamental Science and Innovation

CERN's scientific programme aims at novel and first scientific insight at the energy frontier in particle and nuclear physics. To conceive, construct and exploit the instruments required for this goal, scientists, engineers and technicians at CERN and in collaborating institutes have developed or used as first users new, innovative technologies and combinations of such technologies.

Performances of particle and nuclear physics devices such as accelerators and experimental apparatus have had exponential increases in relevant performance figures over many decades, realised in an environment of, at best, constant funding in the scientific community.

In a number of cases "breaking" technologies and methods have emerged and continue to emerge with considerable interest for other sciences, industry and society. The pervasive technologies having brought most of the exponential progress of the science have been the Information and Communication Technologies.

Systematic exploitation of ICT in particle physics makes it an e-science "par excellence". CERN, with its initiatives LCG, EGEE and others, wants to contribute to the process from scientific insights and applications to novel industrial uses and innovations and to provide for the ingredients of a basic ICT infrastructure enabling such uses.

Fundamental science can act as a long range, medium to long term force towards innovation. Dialog with other sciences and industries is an important means to make the process happen.

An even better way is to find common projects of interest to both industries and the scientists. In some European countries scientific institutions and universities have a triple purpose, namely to do front end science at a global scale, to train competent people and to provide for innovation from the science.

This is certainly true for CERN.

In any case, it is useful for industry to follow closely the progress of (fundamental-) sciences. They occasionally provide for a glimpse into the future of technology and areas of innovation and allow to base strategic decisions of industries on some evidence.

