



HAL
open science

A worldwide grid infrastructure for the four LHC experiments

F. Malek, F. Hernandez, François Chollet

► **To cite this version:**

F. Malek, F. Hernandez, François Chollet. A worldwide grid infrastructure for the four LHC experiments. 2010, 6 p. in2p3-00485291

HAL Id: in2p3-00485291

<https://hal.in2p3.fr/in2p3-00485291>

Submitted on 25 May 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A worldwide grid infrastructure for the four LHC experiments

LCG-France Management: F. Malek, F. Hernandez, F. Chollet

Introduction

The LHC is the Large Hadron Collider, located at CERN (European organization for particle physics) near Geneva. It is the most powerful accelerator in the world in terms of energy, 14 TeV foreseen in the centre of mass of the collision of both hadrons (proton-proton).

The particles are accelerated at a speed close to that of the light and they circulate in a ring of 27 km of circumference buried 100 m underground. The LHC and the associated experimental equipments were conceived and built by international collaborations. It has delivered its first collisions at low energy before 2009 Christmas and it is now operating successfully at a higher energy (7 TeV at c.m.), see **Fig. 1**.




Fig.1: An event at 7 TeV collisions within ATLAS experiment, April 2010. This is a W candidate decaying to a muon and a neutrino

It has required titanic works of infrastructure. It is also a human and intellectual adventure, both unique in the history of science. The analysis of the enormous quantity of data recorded by the four experiments exploit a new worldwide infrastructure, the Grid, comparable in CPU power to the infrastructure of companies such as Google targeted towards an efficient transfer and storage of the data.

The state of the art: WLCG

At the nominal conditions of functioning of the LHC, i.e. 14 TeV proton-proton collisions with luminosity of around 10^{34} particles per second per cm^2 , the experiments will record every year some 10^{10} collisions, a number comparable to the amount of stars in the Milky Way. Data handling represents a real IT challenge, in terms of data flow (of the order of Gigabytes per second) as well as of volume (several tens of Petabytes every year). At any time, thousands of researchers from any part of the world will look for computing resources and the associated storage capacity to be able to analyze the data. To face this challenge, the solution of the distributed calculation, or Grid calculation, was imperative. It was implemented by the collaboration WLCG (World wide LHC Computing Grid).

The WLCG technical Design Report [1] was issued in June 2005, describing the grid infrastructure, the common tools to be developed and the computing models for each of the four LHC experiments. The computing centres providing resources for the Worldwide LHC Computing Grid are also connected to other grids, in particular Enabling Grids for E-science (EGEE) in Europe and Open Science Grid (OSG) in the United States, but also several national and regional grid structures such as GridPP in the UK, INFN Grid in Italy and NorduGrid in the Nordic countries. A Memorandum of Understanding [2] was established in October 2005 between the WLCG collaboration and the participating nations and funding agencies. This MoU guarantees the resources, the quality of services and looks 5-year forward for the computing resource planning. The quality of services includes a guarantee of operations 24 hours a day and 7 days a week with intervention to services essential to the running of a centre within a time laps of 4 hours. For any site in the first level of the hierarchy (also known as tier-1s), the target reliability is 98%. The Worldwide LHC Computing Grid combines the computing resources of more than 100,000 processors from 150 institutions in 33 countries, producing a massive distributed supercomputer that will provide more than 7000 physicists around the world with near real-time access to LHC data and the computing capability to process them. **Fig. 2** shows a snapshot of the WLCG real time monitoring where the worldwide computing centres are processing data and exchanging information.




Fig.2. The real time GridPP monitoring showing the WLCG computing facilities processing jobs and transferring data.

The LHC will produce around 15 petabytes (15 million gigabytes) of data every year for ten to fifteen years. This is enough to fill 3,000,000 DVDs every year. Viewing 3,000,000 DVDs would take around 500 years. If LHC data were to be burned to CD, a tower of CDs around 20 kilometres high would be created within a year. The WLCG infrastructure is based on three “tiers” (Tier0, Tier1, Tier2) and 33 countries are formally involved. The CERN facility (Tier0) is linked to other major national computing centres (Tier1s) using 10 Gigabit per second optical wide area links, as well as to the general global education and research network infrastructure. The four LHC experiments computing models rely on the distributed computing facilities sketched on **Fig.3** which share the responsibilities of the needed computing services: raw data archiving, data distribution, simulated data production, data

analysis More details can be found in the computing models described in the technical design reports for each of the experiments [3].

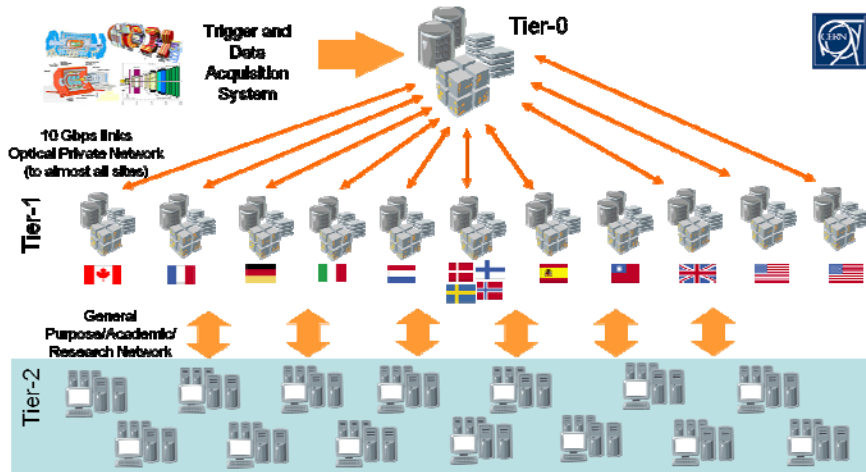


Fig. 3: The computing model

The French contribution to this effort:

Fig. 4 shows the French infrastructure which comprises a Tier1 located in Lyon at the IN2P3 computing centre (CC-IN2P3), and the Tier 2/Tier 3 located in Annecy, Clermont-Ferrand, Grenoble, Ile-de-France, Lyon, Marseille, Nantes, and Strasbourg. This infrastructure and the resources pledged to WLCG are managed through LCG-France project [4].

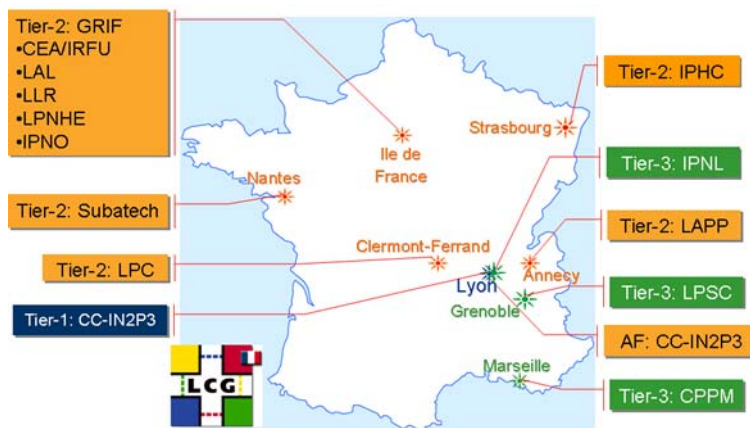


Fig. 4 The French infrastructure operating centres in 2010

The target contribution of the French sites to the worldwide global effort is close to 10%. Fig. 5 shows the actual CPU contribution per country in 2009.

Also, to meet WLCG requirements, reliability and availability tests are performed regularly. **Fig. 6** shows the evolution of those metrics for the French Tier1 since 2006. Actually, to meet the 96% availability means for the site to stay unavailable less than 14 days over the year, including the unavoidable scheduled quarterly shutdown periods for maintenance purposes. The target reliability for the Tier2s sites is 95%. **Fig. 7** shows the monthly reliability scores measured at the French sites in 2009. It is excellent for the majority of sites.

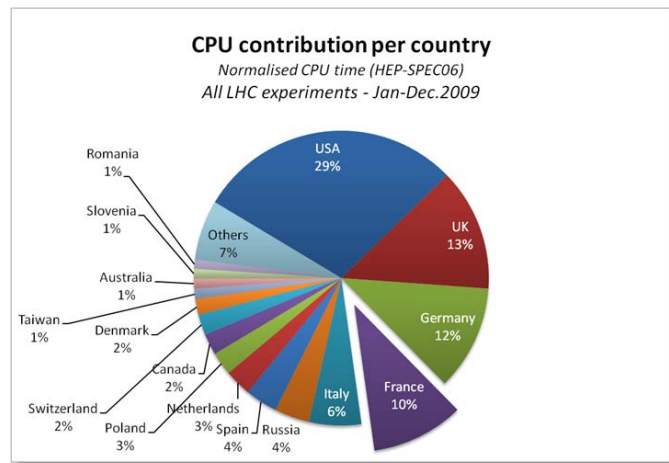


Fig. 5: LCG-France CPU contribution to the worldwide symphony in 2009.




Fig. 6: Reliability and availability of the French Tier1 site since 2006.

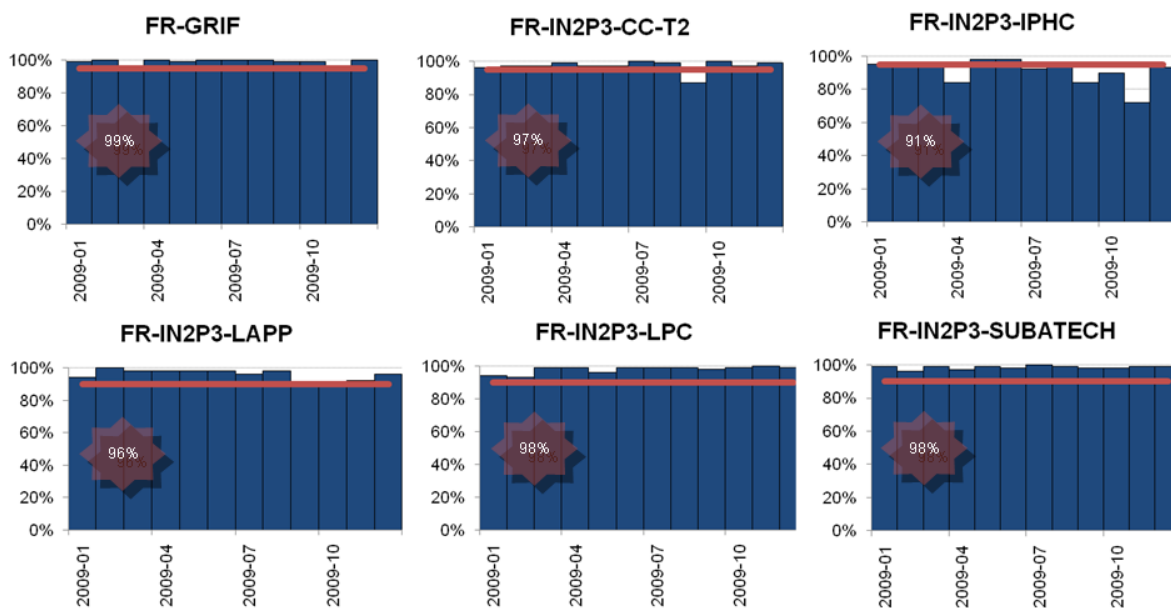


Fig. 7: Reliability of the French Tier2 sites since 2009.

As more data is gathered from collisions inside the LHC, there will be increasing demand upon the storage and processing capacity of the Worldwide LHC Computing Grid. These challenges will include the increase of available resources in response to planned upgrades to the LHC accelerator, as well as to the increasing data requirements of the four LHC experiments. This increase is already observed, as sketched on **Fig. 8** for the French Tier1 site. As of today, it has reached 3.5 PBytes and is planned to double fall 2010.

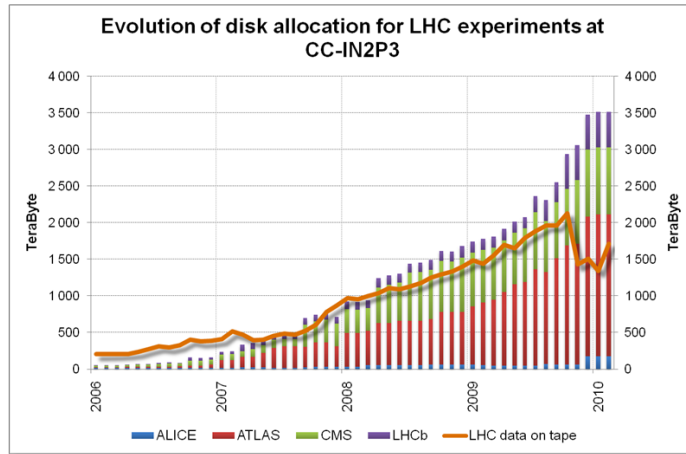


Fig. 8: The evolution of disk allocation and storage for the LHC experiments since 2006 at the French tier-1.

Also interesting to observe is the data processing activity in France, sketched on **Fig. 9** which shows the number of jobs processed since 2006. The amount is reaching 1.5 Million jobs per month in 2010 for the four LHC experiments, the ATLAS experiment being the big consumer.

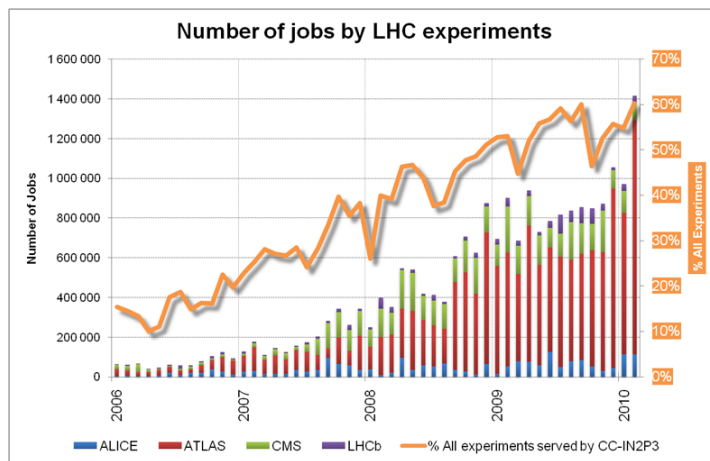


Fig. 9: Number of jobs processed since 2006 at the LCG-France tier-1.

Last but not least, the French infrastructure is very well performing and very much appreciated worldwide thanks to the support of the funding agencies (CNRS/IN2P3 and CEA/Irfu) and to the appropriate budget given to the sites. It has to be noticed that, not less

than 5 Million Euros are necessary every year for equipment and operations of the LCG-France infrastructure, not including salaries.

Conclusion

With the start-up of the LHC, we are about to enter the operational phase of the WLCG infrastructure. After intensive successful testing in the last few years, we will see what it delivers in real circumstances. Priority is now to make the grid more useable for the individual scientist and in particular for data analysis.

The WLCG applications support and service meet with satisfaction the requirements and the baseline services are in production. There is, as expected, a continuously increasing capacity and workload and the general site reliability is improving. It has already reached a quality level of 98% for most of the Tier1s, the target value fixed by the WLCG MoU.

Besides this, adoption of new technologies is important. Continuing to optimize the applications to fully exploit multi-core processors, where cores increase beyond two or four or eight processors to 16- or 32-core processors is essential. In the medium term, some technical work will also focus on improving the performance and scalability of access to data, especially for analysis use cases. And finally, working with limitations in terms of the cooling and power requirements of large data centres, is an ongoing issue shared by large data centres all over the world.

[1] CERN-LHCC-2005-024, see also: <http://lcg.web.cern.ch/LCG/tdr/>.

[2] CERN-C-RRB-2005-01.

[3] ALICE-TDR-012: CERN-LHCC-2005-018;

ATLAS-TDR-017: CERN-LHCC-2005-022;

CMS-TDR-007: CERN-LHCC-2005-023;

LHCb-TDR-11: CERN-LHCC-2005-019.

[4] LCG-France project: <http://lcg.in2p3.fr/>