

LHC Computing Grid

L. Aphecetche, J.-M. Barbet, K. Chawoshi

Le futur accélérateur du CERN produira une quantité telle de données qu'il sera impossible de les traiter sur un seul site ou même dans un seul pays. L'analyse de ces données reposera donc sur une infrastructure de grille informatique répartie dans le monde entier. Le laboratoire Subatech participe à cet effort mondial en fournissant, à l'horizon 2008, 280 kSI2K de puissance de calcul et près de 80 téra-octets de stockage, au sein d'un Tier2 LCG.

The Large Hadron Collider (LHC), currently being built at CERN near Geneva, is the largest scientific instrument on the planet. When it begins operations in 2007, it will produce roughly 15 Petabytes (15 million Gigabytes) of data annually, which thousands of scientists around the world will access and analyse. The mission of the LHC Computing Grid Project (LCG [1]) is to build and maintain a data storage and analysis infrastructure for the entire high-energy physics community that will use the LHC.

Being highly involved in ALICE, one of the four LHC experiments, the Subatech laboratory decided to invest also in LCG, to provide computing power, storage capacity, and human resources.

The LHC Grid is hierarchically organized in computing centres of various computing capacities and responsibilities, denoted as "Tiers". The single Tier 0 will be CERN and will span, for Alice, all computing activities, except simulation: permanent mass storage of raw data and reconstruction and physics analysis. Less than 10 Tiers 1 « big » centres will perform almost the same tasks as Tier 0. Finally, a bigger number of Tiers 2 (T2) will be devoted to Monte Carlo simulations and physics analysis. T2 typically do not have a mass storage system and do not deal directly with the raw data. Subatech is setting up such a T2, building upon our existing (small) computer farm that has been in use since 2000.

This grid activity is mainly in the hands of the IT service (0.8 FTE), overviewed by a physicist (0.2 FTE). The tasks are both software and hardware in nature.

On the hardware side, besides being alert to innovations in the fast-moving computing market, we must both operate and maintain what we already have, and handle our procurements of new hardware (CPUs or disk storage). Already installed and operating on site, we have:

- 20 Pentium IV 3.0 GHz 1 GB RAM
- 20 Pentium IV 3.2 GHz 2 GB RAM
- 5 TB of storage in a single bi-opteron server

The next steps for this year are an extra 13 TB of disks (IBM DS4800) and a doubling of the number of CPUs.

The evolution of the part of our capacity that will be engaged in LCG is depicted in Fig. 1.

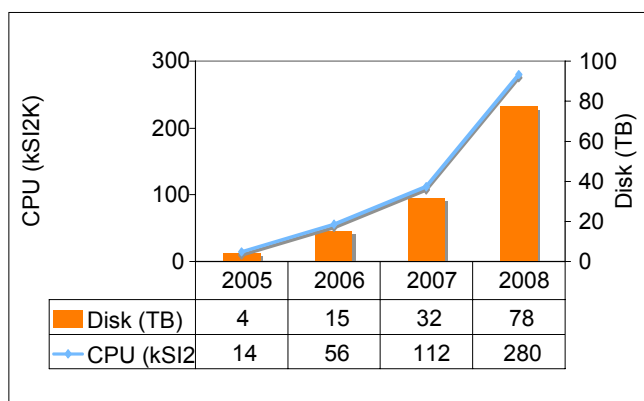


Fig. 1: Planned evolution of the Subatech T2

On the software side, we must install and maintain all the LCG grid middleware, as well as some ALICE specific components, bundled into the AliEn [3] system. All these software packages are under intense worldwide development as we approach the LHC turn-on (envisaged mid-2007), so that the operation of this complex web of software services is still not a “plug and play” task and requires a fair amount of debugging.

We must also ensure that the operation of our T2 has a minimal impact on our local non-LHC users, who use part of this computing power.

The hardware computing grid is a highly distributed environment. Likewise, the people that make the grid work are scattered around the world, so that this activity is not a standalone process. We are part of the LCG-France collaboration [2] and a member of the WLCG collaboration board. We are also the representative of the French T2s in the ALICE collaboration.



- [1] <http://cern.ch/lcg>
 [2] <http://lcg.in2p3.fr>
 [3] <http://alien.cern.ch>

Qu'est-ce que la grille ?

Comme la toile, la Grille est un service permettant de partager quelque chose sur l'internet : l'information quant à la toile, la puissance informatique et la capacité de mémoire quant à la Grille. Mais cette dernière fait beaucoup plus qu'établir une simple communication entre les machines. Son objectif ultime est de transformer le réseau mondial des ordinateurs en une ressource immense et unique de capacité de traitement.

La grille de calcul du LHC

Le LHC produira environ 15 pétaoctets de données par an, soit plus de mille fois la quantité d'information contenue dans tous les livres imprimés sur la planète et près de 10% du volume total d'informations produit par l'homme - y compris les photographies et les images numériques et tout ce que vous pouvez imaginer.

Le défi auquel est confronté le projet LCG ne tient pas seulement à la grande quantité de données que les physiciens devront engranger car pour analyser ces données du LHC, il faudra environ 100 000 de nos ordinateurs individuels actuels !



La seule façon raisonnable d'accéder et de traiter cette grande quantité d'informations semble être la technologie de grille.

LCG est actuellement la plus grosse grille au monde avec 180 sites, 35000 processeurs et 9 pétaoctets de stockage.