CERN COMPUTER NEWSLETTER

Contents

IOP Publishing

PSS group ensures efficient use of Grid
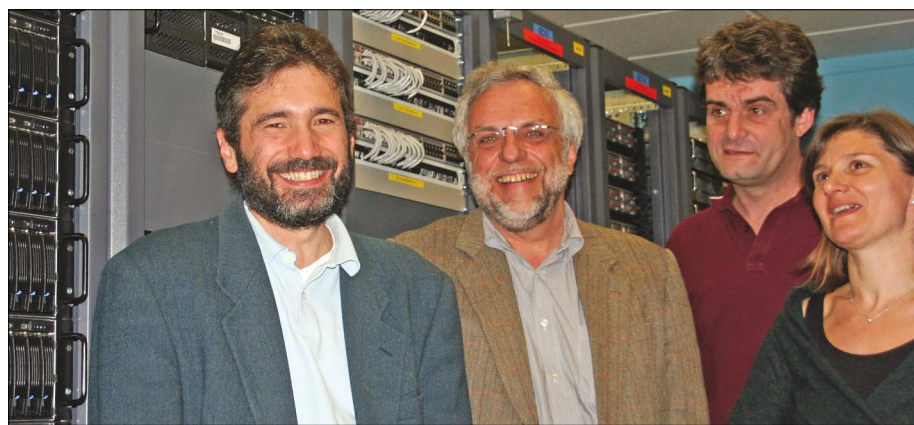
PSS group's Massimo Lamanna (Arda/EIS project leader), Jürgen Knobloch (PSS leader), Dirk Düllmann (LCG database project leader) and Maria Girone (physics database services).

This interview with Jürgen Knobloch, leader of the Physics Service Support (PSS) group, continues a series this year in which *CNL* is focusing on groups in the IT department that provide key support to the LHC Computing Grid (LCG). PSS is a relatively new group that was formed in a reorganization of department activities in November 2005. The group, of about 50 people, is supporting the rapidly evolving needs of the physics community as a production Grid service comes online and preparations to analyse the first LHC data intensify. PSS has two main sections, one devoted to databases for physics, the other to experiment support and distributed analysis.

What is PSS doing regarding databases to prepare for the start-up of the LHC?

We run the physics databases for the CERN experiments. For LHC experiments we store calibration and alignment data as well as Grid file catalogues. The ORACLE servers, currently more than 200 CPUs, receive the constants from the data acquisition systems and other sources to be ready for the first-pass reconstruction at CERN.

We also support the replication of the databases from CERN (the Tier-0 centre of the LCG) to the Tier-1 centres. This is done for ATLAS and LHCb using Oracle Streams, and for CMS through a web-caching mechanism called FroNTier, which is based on Squid servers. This technology ensures that changed records are transmitted to the other centres so that the Tier-1 databases are synchronized with those in Tier-0. The physics databases at CERN play a central role in the computing strategy of the experiments, and they are designed and operated to be reliable.

One of the challenges we've faced is that this technology hasn't been tried on such a scale before, so we have discovered certain limits. We have made significant improvements through the CERN openlab partnership with Oracle, which is sponsoring two people in PSS to work on this. In particular, in recent months we have overcome a strong latency effect. When shipping data from CERN to places like Taipei this made a difference of an order of magnitude in the rate at which we could refresh distributed databases.

This whole system of distributed databases will be probed soon, when the experiments start their calibration challenges. Until now most Grid work has not relied on local copies of databases, but from April all the experiments will test their ability to use up-to-date calibration constants for the detectors when reconstructing data. We anticipate that

Editorial

this will be a challenging time for us.

We support the persistency framework for LCG, which has been developed in close collaboration with all four LHC experiments. This framework provides a software layer that decouples the user code from the features of any particular database implementation. The persistency framework project focused initially on the development of POOL, a hybrid store based on object streaming into ROOT files and metadata storage into databases. More recently the scope of the project has been extended to provide a generic database access layer (CORAL) and a specialized component for storing and looking up conditions data (COOL).

How does PSS help physicists use the Grid?
Our experiment support and distributed analysis section is involved in certifying new releases of the Grid middleware, gLite. We are also involved in getting the experiment data management systems to work. Furthermore, we support job submission facilities such as GANGA, which was developed by LHCb and has been taken up by ATLAS. We try to foster commonality between experiments, although this is often challenging given their tendency to want to be independent of each other.

Another example of getting experiments to agree on common support is the dashboard used to monitor how jobs behave on the Grid. The dashboard was developed in the ARDA project and was originally used by CMS. It has now been adopted by all the experiments. It's useful because it helps users to understand job failures. Inefficiencies on the Grid can be identified and remedied with this tool.

The role of PSS is to maintain the dashboard and adapt it to the evolving needs of the users. For example, you can now use it to measure the efficiency of individual Grid centres. Each experiment may have problems running jobs on a different set of centres, and the dashboard helps them know where to look for improvements. Since March we have been reporting Grid efficiency statistics regularly to the LCG management board. LCG is now running jobs at nearly 90% reliability at most sites. As we approach LHC start-up the goal is to steadily increase this figure.

As well as high energy physics, what other scientific initiatives do you support?
We receive a lot of support from the Enabling Grids for E-sciencE project for our work with the LHC experiments. In exchange, we are expected to devote some effort to helping other scientific communities adopt Grid technology. This provides our group with useful experience. For example, we have helped to gridify satellite imagery applications for UNOSAT, a United Nations initiative that provides the humanitarian community with access to satellite imagery for use in crises such as earthquakes and tsunamis.

We also supported a series of large-scale data-processing activities carried out by the International Telecommunications Union last year, during a five-week conference to establish a new frequency plan for digital broadcasting in Europe, Africa, the Arab states and the former USSR. By using a few hundred computers on the Grid, the most demanding analysis step could be reduced during the conference from about four hours on a local cluster to less than one hour on the Grid. At the moment we have a doctoral student using Grids and the Geant4 software package to carry out simulation tools for hadron therapy.

Looking beyond the LHC start-up, what challenges lie ahead in your view?
We must face the fact that on the database side the experiments themselves do not know exactly what their requirements will be. Everything will evolve with experience, once real data starts to flow.

For example, we do not know how fast the "constants" that describe the detectors will vary, because there will always be time dependencies of the calibration. How often you need to update constants remains to be seen. Indeed, the number of constants needed is still not known. For example, when we started to work on the Large Electron Positron collider experiments physicists thought they would only need a few constants. But in the end it turned out that each detector element needed a separate calibration.

So we have to be flexible when predicting the requirements for detectors. We have now installed about 300 TB raw capacity for the databases, which will be doubled by the time the LHC starts up. We hope this will get us through most of 2008. Beyond that we must be prepared to develop new solutions.

We can also anticipate new applications for databases, which will prove challenging. For example, in ATLAS some physicists are discussing storing all the metadata, the so-called tag data, in Oracle databases. This would be a huge amount of data and we probably wouldn't be able to distribute it to the Tier-1 centres using the streaming techniques we have at the moment, so we would need to study new approaches.

For me the crux of this business is data management. The experiments will need to find the right balance between flexibility and pragmatism. Many researchers would like to see the Grid like a single supercomputer but the Grid cannot provide this sort of functionality. There are limitations due to latency, network throughput and storage capacity. Existing Grid tools will need to be tuned to the needs of the physicists. And the physicists will need to be tuned to the realities of the Grid.

Computing articles featured in this month's CERN Courier

The articles listed below appear in the April issue of *CERN Courier*. Full-text articles and the rest of the issue's contents are available at www.cerncourier.com.

Computing News

● **Worldwide Grid collaboration gets ready for LHC start-up**
Week-long workshop at CERN prepares WLCG members for launch of LHC.
● **Worldwide Grids create worldwide science**
The US science association, AAAS, discusses distributed computing.
● **WISDOM ends second round in the battle against malaria**
Grid analyses 140 million compounds.
● **W3C celebrates 10 years with style**
The anniversary of CSS technology.
● **Commercial quantum chip goes on show**
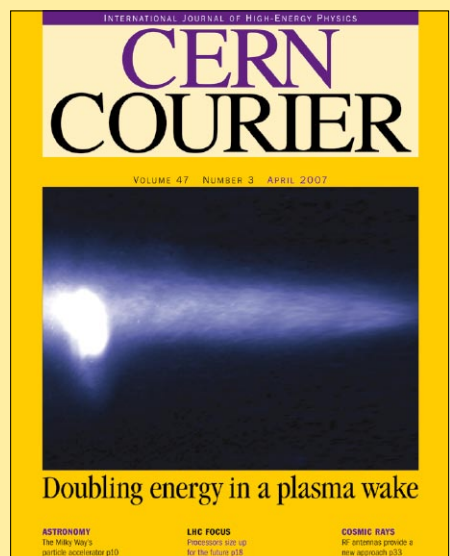D-Wave demonstrates the Orion chip.
● **Chip may boost particle-physics Grid power**
Quad-core processor increases speeds.
● **How to avoid a false sense of security**
Advice on protecting PCs from infection.

Calendar of events

CERN COURIER

INTERNATIONAL JOURNAL OF HIGH-ENERGY PHYSICS

VOLUME 47 · NUMBER 3 · APRIL 2007

Doubling energy in a plasma wake

ASTRONOMY
The Milky Way's particle accelerator p10

LHC FOCUS
Processors size up for the future p18

COSMIC RAYS
RF antennas provide a new approach p33

Feature article
● **Processors size up for physics at the LHC**
Software challenges in the LHC era and the evolution of PC processors.