1966 40 YEARS OF CNL 2006

# CERN COMPUTER NEWSLETTER

Volume 41, issue 4 September–October 2006

## Contents



**Editors** Nicole Crémel and Hannelore Hämmerle, CERN IT Department, 1211 Geneva 23, Switzerland. E-mail: cnl.editor@cern.ch. Fax: +41 (22) 7668500. Web: cerncourier.com/articles/cnl.



**Produced for CERN by Institute of Physics Publishing**
Dirac House, Temple Back, Bristol BS1 6BE, UK. Tel: +44 (0)117 929 7481. E-mail: jo.nicholas@iop.org. Fax: +44 (0)117 920 0733. Web: iop.org.

**Published by CERN IT Department**



# Common controls help experiments

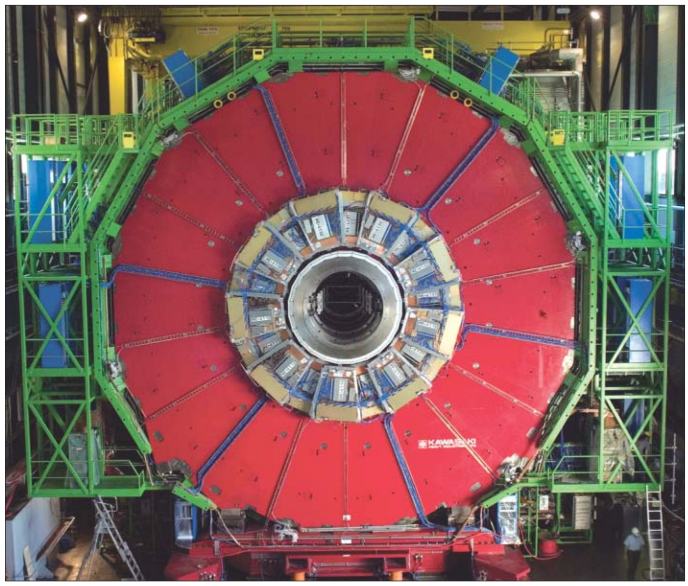

*The Controls Group supports the common controls software that is used by all of the LHC experiments. The CMS is pictured above.*

The IT department's Controls (CO) Group supports much of the software that the control systems for the Large Hadron Collider (LHC) experiments rely on. In the count down to the LHC start-up, the leader of the CO Group, David Myers, and his deputy, Wayne Salter, spoke to *CNL* about the group's activities and future challenges.

**Can you outline the CO Group's history and activities, and tell us how much effort these require?**

DM: In the era of the Large Electron Positron (LEP) collider no central team provided support to the experiments for control systems software, and this led to duplication and problems of long-term support. Our group was originally set up in the Electronics and Computing for Physics (ECP) division in 1993 to improve this situation.

In December 1997 I was asked to set up the Joint COntrols Project (JCOP) to formalize discussions on this subject between representatives of CERN IT and the LHC experiments. At the highest level of management of the experiments there was general agreement on the need to do things in common, and the CO Group's main role has since been to implement the JCOP's proposals. The basic idea was that controls software was no longer rocket science so we could use commercial solutions rather than start from scratch.

For example, Supervisory Control And Data Acquisition (SCADA) is a well established concept in industry, with specialized software vendors offering solutions that can be adapted to the high-energy physics (HEP) community's needs. At first the experiments needed convincing that industrial solutions could be used, but through JCOP we have been able to define a range of areas

where joint efforts make sense, and we have achieved much standardization – more than people anticipated at the outset.

Today, for SCADA, we support PVSS II (*Prozessvisualisierungs- und Steuerungssystem* – process visualization and control system) developed by the Austrian company ETM. We also maintain a centre of excellence in industrial and custom controls hardware and software, including field buses and middleware, as well as maintaining test labs for evaluating actual hardware. Further, we provide support for LabVIEW – a dataflow visual programming environment from National Instruments (NI) – as well as other NI hardware and software products. In some cases we actually develop complete software systems. CO now has 16 staff, with another 10 associates and students.

**How does your support of products like PVSS and LabVIEW help the experiments in practice?**

WS: PVSS is generic – it knows nothing about any specific piece of hardware. We provide much customization in the form of the JCOP Framework. As an analogy, consider Microsoft Word, where you can make any document from scratch but it is more efficient to have templates for standard documents like letters. In the same way, our role has been to create a set of PVSS templates that target HEP-specific needs, such as high-voltage control. With these templates the experiments can build an application-specific layer quickly, saving weeks of work each time they create or modify a control system.

With LabVIEW, as for PVSS, our support includes licence handling, software distribution, testing new versions and debugging. An important example of a LabVIEW application is the facility in SM18 for testing all of the cryogenic magnets. In some cases CO support extends to expert help for individual projects, such as data acquisition and analysis for electron beam welding, and a portable system to analyse radiation sensor connectivity and noise.

It is worth emphasizing that although the CO Group's mandate is primarily to support the LHC experiments, for PVSS and LabVIEW we provide CERN-wide support; for example, for the accelerators and other experiments. We support PVSS for LHC users at 100 partner institutes in 26 countries around the world, and we have many hundreds of PVSS and LabVIEW users in total. Thanks to the Controls Board, other lab-wide controls support has been similarly centralized in other departments.

**What sort of software systems has the CO Group been involved in developing?**

DM: One example is the Gas Control System (GCS). The LHC subdetectors have 23 different gas systems; for example, wire chambers require control and monitoring of gas flow, gas temperatures, pump speeds and so on. In the past these systems were developed independently and required several people per system. We are generating the software for these 23 systems centrally with about five full-time equivalents (FTEs) – a considerable saving.

In practice we create customized instances of a generic model of a gas system based on the information from our colleagues in the Physics Department. From the instances, we then generate all of the software for the Programmable Logic Controllers (PLCs) as well as the PVSS supervisory code. We have completed versions of nine systems so far.

Another example is the Detector Safety System (DSS). Given the huge cost of the equipment involved in the LHC experiments, it is vital to monitor temperature, humidity and other parameters continuously to detect any anomalies and switch off power if equipment is at risk. We have made five DSSs, two for CMS and one for each of the other LHC experiments. Again, this approach has avoided one-off developments of similar systems and has been done with two FTEs.

**What challenges do you foresee for the launch of the LHC and beyond?**

WS: There's a big difference between running a small set-up in a lab with one module and operating a full detector. And there is yet a further step in complexity when you want to operate multiple detectors alongside each other. Thus we will surely run into integration issues over the coming months. People simply haven't dealt with such complex control systems before. Previous control systems for LEP had only a small number of computers, whereas for the LHC experiments there will be a hundred or more.

DM: Another issue is that PLCs, like PCs, are nowadays commonly connected to the Internet and so are under attack by viruses and the like. Losing data on your PC due to a virus can be very upsetting, but losing control of parts of an LHC experiment due to a virus could potentially cause major physical damage, so this is a matter we are taking very seriously.

One challenge is how to provide security patches to PCs that are performing critical functions and therefore can't be turned off. A fundamental issue has been to convince manufacturers of PLCs that there is indeed a security threat to their equipment. Mandated by the CERN management, we have launched the Computing and Network Infrastructure for Controls (CNIC) working group to develop a security policy for the lab in this area.

Beyond the immediate start-up issues, experience shows that as the physics community start to understand their new detectors better, they will inevitably request many improvements to get the most out of the equipment. Thus, we will surely continue development work over the coming few years.

**Do you think the lessons learned from the CO Group's efforts could be relevant elsewhere?**

JCOP has been a very successful example of streamlining resources. Of course, having the experiments and various accelerator systems rely on centralized support demands careful planning to ensure that qualified personnel are always available, but the resulting savings are enormous. It is encouraging to note that we are already being approached by engineers at ITER, the experimental nuclear fusion reactor being built in Cadarache, who are keen to learn from our experience how to run their controls more efficiently.

# Computing featured in this month's CERN Courier

The articles listed below appear in the September 2006 issue of *CERN Courier*. Full-text articles and the rest of the issue's contents are available at www.cerncourier.com.


## Computing News

- **Grid helps achieve international digital-broadcasting agreement**
- **Bristol gets new computing power**
- **Biologists use EGEE Grid to find drug compounds to battle avian flu**
- **W3C announces Web authoring for mobiles and a new Web-services recommendation**
- **The WWW springs into action**
- **Info.cern.ch is live once more**
- **DIANE creates efficient scheduling on the Grid**
- **GEANT2 launches screensaver to celebrate success**

## Calendar of events

## Feature article

**Nordic Grid activities strengthen the WLCG**

# Computing at CERN in the 1990s

In the first three issues of *CNL* this year we looked at the history of computing at CERN up to the 1980s. In the nineties the most important event in computing was probably the birth of the World Wide Web (WWW). This issue will focus on articles related to this revolution. The first Web version of *CNL* appeared in October–December 1993, and you can read it at http://cern.ch/cnlart/214.

We have found Tim Berners-Lee's first articles for *CNL*. This shows how important it is to write down information that can be archived and retrieved later. We hope that this will encourage people to submit articles to *CNL*, especially when working on new projects, to describe and document the various phases.

For a detailed explanation of how the WWW project started and why, check out Tim Berners-Lee's article 'HyperText and CERN', dated March 1989 and May 1990 (www.w3.org/Administration/HTandCERN.txt).

*Tim Berners-Lee explains how the Web came to be. He is pictured here at the Internet, Web, What's Next? conference held at CERN in June 1998.*

**A project for hypertext and browsable documents at CERN**

"Depending on its complexity and distribution, a document will be presented to the users in one of the following forms:
- on paper (major manuals, or document widely distributed outside CERN);
- as plain text for simple terminals (we hope all documents);
- as a sophisticated version for workstations – these versions will be developed in the context of the hypertext project (currently being launched between ECP and CN).

The major long-term aim will be to make machine readable versions browsable from outside CERN." (July–October 1990)

**World Wide Web: online information for everyone**

"A world of information is now available online from any computer platform... We summarize the information currently sourced at CERN, and we introduce the World Wide Web (W3) program, which allows you to browse and search all the data in a simple and consistent manner.


*Berners-Lee developed the first WWW server, multimedia browser and Web editor on this machine in 1990. It is now on display at CERN's Microcosm science centre.*

Details of all the information, and of how to acquire and install the W3 software on your own machine, are available by typing telnet info.cern.ch...

Information... is usually stored on one platform (the machine, the format and the application that controls the information). The platform has peculiarities which make it different from other platforms. However most machines are connected by networks, and most systems have (terminal) emulators through which they give access of some level to other platforms. Therefore, there are three ways to access information:
- go to the platform the data is on, log into it, use the application specific for these data;
- connect to the platform, from another platform, and use it from there; or
- use a platform-independent application.

In the third case, you can navigate through the information without having to know details depending on a platform. In the following section we will show how to invoke the platform-independent application: W3. To date, only W3 can give you platform-independent access...

The browser described here is the simplest interface to the World Wide Web, which was designed to run on any dumb terminal. A hypertext browser/editor is available under NeXTStep. More powerful interfaces are under development, including Macintosh, X-Windows, VM full-screen and emacs...

T Berners-Lee, R Cailliau, J-F Groff, B Pollermann." (October–December 1991)

**W3: spring releases**

"The World Wide Web bounds into spring with some exciting software releases from outside CERN. A few more servers pop up, and we look at the Gopher and WAIS part of the Web...

Statistics we take on the access to our server show an exponential increase, with a time doubling about every two months, and currently running at around 500 document fetches a day... many of these are from people telnetting to info.cern.ch...

***X11 browsers: Erwise*** For those who have waited and asked so many times, here are not one but two client products which run under X windows... The Erwise browser was developed at Helsinki Technical University... It is a multi-window, multi-font hypertext browser which uses the W3 common access library and so has full W3 functionality...

***X11 browsers: Viola*** The Viola browser by Pei Wei of Berkeley is another interesting product... Pei's solution to the multiwindows/single window dilemma is to do everything in one window by default, with a 'clone' button which duplicates the current window...

***W3 people at CERN*** Jean-François Groff, who has been a major mover in the W3 architecture, and has been invaluable in keeping the show on the road, came to an end of his contract at the end of March." (March–April 1992)

**New version of Mosaic available**

"Mosaic is an X11-based application to browse information stored in the World Wide Web system. It has been developed at NCSA and provides at present the best user interface available on Unix workstations.

We have released into the Public Area the Unix, VMS and Macintosh versions of Mosaic 2.2.

In addition Mosaic is available on PC/Novell as provided by Robert Cailliau/ECP...

PCWeb is in fact the SUN version of Mosaic, displaying on PCs which have access to the XVision product. There is an icon called World Wide Web in the Comms folder inside Program Manager. The inconvenience of this version... is that due to the way it has been set up, Mosaic will NOT remember the documents you have 'visited'." (January–March 1994)

**Using CERN's phone book from W3/Mosaic**

"Two kinds of documents in the CERN-wide Web are of interest...

1. documents informing people about software, projects services etc.; and

2. documents informing people about people, especially their e-mail addresses and phone numbers.

In the first case a reference to the CERN phone book is needed for the signature which indicates the author or person responsible for the document or service... In both cases it is useful to have a direct pointer to a particular person... This identifier, called ccid ( Computer Centre IDentifier), was introduced some years ago and can be used now also in W3/Mosaic." (April–June 1994)

# Real-time status display for IT

Most of us at CERN regularly use one or more of the computing services provided by the IT department. Until recently, if we experienced problems we could find out if a particular service was working properly by checking the IT status display (TVScreen).

However, the distributed nature of today's computing services meant that the TVScreen was often misleading. The TVScreen status for a service depended on the number of alarms for the machines supporting the service, and not on the status of the service itself. This led to at least two problems:

- Many of today's services, such as the fax gateway, run on multiple machines for redundancy and so the service can be working perfectly even if one of the machines has a problem.
- Large services like AFS could be almost fully available, but a problem on one machine might affect a particular group of users, such as ATLAS.

## Introducing Service Level Status

To solve these problems and to improve the user visibility of the real status of IT department services, a Service Level Status (SLS) Overview project was launched early in 2006. The aim was to provide a Web-based tool that dynamically shows service availability, basic information and statistics about IT services, as well as their dependencies and inter-relationships.

The SLS Overview display, at http://sls.cern.ch/sls, went live at the beginning of July. It provides a comprehensive and easy-to-interpret display of the status of the services provided by the IT department, as measured in a way that is relevant to users.

Although targeted at end-users, the display provides a high-level view of service status that is useful for IT department managers as well as other service providers. SLS also proved to be useful after the recent power cuts and cooling problems that affected IT services, because it provided a simple overview of which services had been restored and which needed further attention.

As well as being a Web page in its own right, SLS has replaced the TVScreen snapshot on the real-time status page (at http://cern.ch/it-support-servicestatus/default-dynamic.asp) that users can access from the IT Service Status Board (SSB, at http://cern.ch/it-support-servicestatus or http://cern.ch/SSB). In addition to the SLS overview, the SSB provides considerably more user-oriented information, such as details of forthcoming service changes and scheduled interruptions. The SSB is edited daily by the IT manager on duty (mod@cern.ch) with information that has been provided by the IT service managers.

## Extensive coverage

The status of more than 140 computing services, provided by the IT department or in collaboration with other departments (such as the Engineering Data Management Service, EDMS), is displayed by SLS. Services covered include:

- administrative applications (such as EDH and ERT);
- Windows, mail and Web services (such as MMM, Terminal Services and DFS);
- documents and collaborative services (CDS, Indico, CRBS);
- services for physics (LXPLUS, LXBATCH, LXBUILD, Backup, AFS, Tape Service, Kerberos, Linux);
- LCG Tier-0 and Tier-1 sites;
- services for developers and engineers (such as EDMS, CVS, J2EE Public Service and TWiki);
- networking infrastructure;
- multiple databases;
- and many others (such as Remedy, and SLS itself).

To make sense of the status information for so many services, SLS groups them by functionality into a logical structure and provides the ability to drill down to find more detailed status information. The default SLS entry page is shown in figure 1 and users can drill down as necessary to find, for example, the details of the different services for physics users as shown in figure 2.

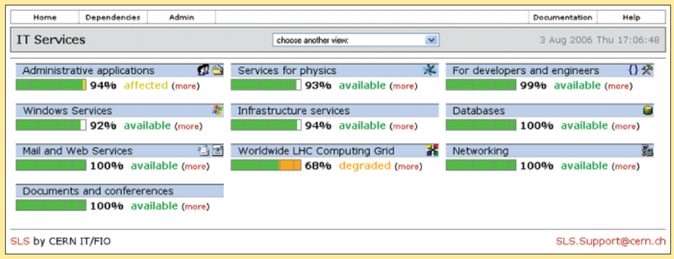


*Fig. 1: The SLS entry page provides a logical structure of services.*

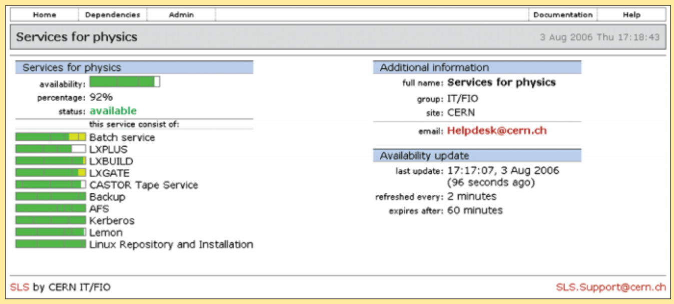


*Fig. 2: The list of services outlined in the Services for physics group.*

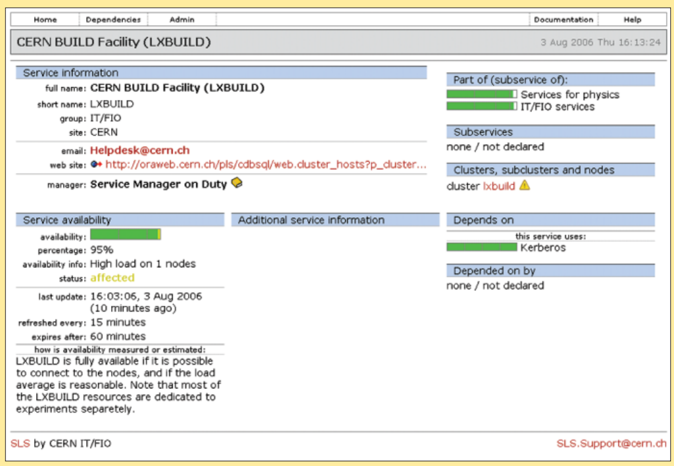


*Fig. 3: Details and availability information for the LXBUILD service.*

Other high-level views can be selected in the drop-down list at the top of the main SLS page. At present four alternative views are available, showing:

- all IT services on one page;
- the status of services provided by the different IT groups;
- the status of Grid services at sites in the Worldwide LHC Computing Grid;
- the status of services for different virtual organizations (or communities). At present this shows the state of services for the different LHC experiments, indicating, perhaps, that ATLAS users are affected by an AFS problem while services are fine for everybody else.

## Interpreting the SLS displays

The "service availability" as shown by SLS indicates the extent to which a given service is accessible and performing as expected. Availability is represented as a percentage. Two values have obvious and common meanings: 100% means that the service is fully available and 0% indicates that the service is completely down.

As the availability values are calculated by the services themselves, intermediate values are relevant to the way in which different services are used. For the batch service, the availability is simply the number

of nodes that are up and running, but services can take many factors into account when calculating the availability. For example, the LXBUILD service for software compilations takes into account the load on the nodes as well as their availability. Figure 3 shows that even with all nodes fully working, a high load on one node reduces the displayed availability to 95%.

In addition to the simple percentage indication of service availability, SLS also shows a status level for the services. Four status levels are available: fully available, affected, degraded and not available. These status levels are indicated by the colours green, yellow, orange and red respectively. Like the percentage availability, the status level is decided in a way that is relevant to the service. For the fax gateway, either of the two redundant nodes is sufficient to provide a full service, so a 50% availability will be shown as a fully available service. On the other hand, for more complex and critical services such as CERN's administrative applications, 95% availability means the service will be shown as being affected and an availability of 50% will lead to the service being shown as not available.

Finally, figure 3 also shows the way in which SLS can help people to understand why services may be unavailable, because it shows that LXBUILD depends on the Kerberos service. If LXBUILD is unavailable and the Kerberos display on the right is red then the real problem is with Kerberos, and there is no point in contacting the helpdesk (perhaps by clicking on the mailto: link at the top of the page) to complain about the LXBUILD service.

In the next few months we will add a display to provide data of historical availability and also displays of key performance indicators such as the fraction of the target batch quota that we deliver to the different experiments.

Add http://cern.ch/sls to your bookmarks and look at the SLS display if ever you think there may be a problem with one of our services. Or keep a display open permanently – you may be surprised to see that services run smoothly most of the time.

**Sebastian Lopienski, IT/FIO**

# Computer Centre upgrade is reaching the end

"Plus ça change, plus c'est le même chose." Yet again one of the most visible aspects of the ongoing upgrade to the Computer Centre is the lack of parking spaces. Back in 2003 part of the car park on the CERN restaurant 2 side of the Computer Centre was inaccessible during construction work for a new substation. That work was completed in 2004 and the substation has been in service since last December. Today the car park on the Salève side of the building is inaccessible because we are upgrading the Computer Centre's air-conditioning system.

This work is necessary because the large computer farms for Large Hadron Collider (LHC) offline computing will need up to 2.5 MW of power. Most of this power is turned into heat, which needs to be removed from the machine room.

*Clockwise from top left: Removing the old cooling towers... three new chillers waiting to be installed... going up... and in place on top of B513.*

*Ducts and grills provide cold air directly to computer air intakes.*

*Dedicated pumps in the basement provide targeted air conditioning.*

## Time for change

The air-conditioning system in the Computer Centre was designed in the early 1970s to maintain an even temperature across the machine room, with most of the excess heat from the mainframes being removed with their dedicated water-cooling systems. With today's PC-based compute farms the heat is released directly into the machine room and a constant supply of cold air is needed for cooling.

The most visible impact of this change is the installation of several large ducts to pump cold air under the false floor rather than it trickling down from the ceiling level. Along with the installation of grills in the false floor, this means the cold air is fed directly to the air intakes of the computers rather than mixing with the hot exhaust air first.

A lot of cold air, more than 500 000 $m^3/h$, is needed to cool the PCs, and this in turn requires a lot of chilled water to cool the air from outside. Just as for the electrical distribution, the 1970s-era chillers and water-distribution network needed a major overhaul to cope. Once again, the space available in the building was not sufficient and so three new 1.5 MW chillers have been installed on the roof of B513, replacing the old cooling towers. This solution also has the advantage of eliminating a possible breeding ground for the bacterium that causes Legionnaire's disease.

With the chillers moved to the roof, the cooling centre in the basement of B513 now has space to provide dedicated circuits, each with redundant pumps, for the two machine rooms in the Computer Centre, for the air conditioning needed for the electrical substation, and for the air-conditioning system in B31. This is a great improvement over the previous continuous circuit because it enables work to take place on the air conditioning in the IT auditorium, for example, without affecting the air conditioning in the Computer Centre.

## Up and running

Unfortunately the weather has not been kind to us during this upgrade. The wet weather in March and April delayed the installation of the new chillers and, with the lack of chilled water, the unexpectedly hot weather at the end of March meant many systems in the Computer Centre had to be shut down until temporary chillers could be installed. Various teething troubles with the new installation then led to further problems in May and early June. Fortunately these problems seem to be over and the Computer Centre is now nice and cool again and ready for the next batch of computers to be installed in September.

**Tony Cass, IT/FIO**

# Persistency framework manages LCG databases

Managing data in a distributed and heterogeneous Grid environment is one of the most challenging tasks of Large Hadron Collider (LHC) computing, in terms of both developing the required software and deploying the underlying services.

Relational database management systems (databases) play a central role, especially for conditions and event metadata, because they can provide consistent storage to many concurrent users.

Given the complexity of today's database systems, it is often difficult for users to exploit the underlying services in the most efficient way. However, many of the required optimizations can be delegated to an intermediate software layer if both the main physics use cases and service constraints are taken into account early on in its design.

### LCG persistency framework

The persistency framework for the LHC Computing Grid (LCG), which is being developed jointly between the LHC experiments and the IT/PSS group, aims to provide such a software layer. Its purpose is to decouple the user code from the features of any particular database implementation.

The project started in 2002 in the LCG Applications Area, driven by the requirements of its users (ATLAS, CMS and LHCb experiments). Project priorities are set with the experiment representatives in the LCG Architects Forum, and experiment developers contribute actively to the software implementation. Development is also tightly coupled to service constraints, as the software is developed in close contact with the IT/PSS physics database service team and the LCG Distributed Deployment of Databases (3D) project (led by IT/PSS).

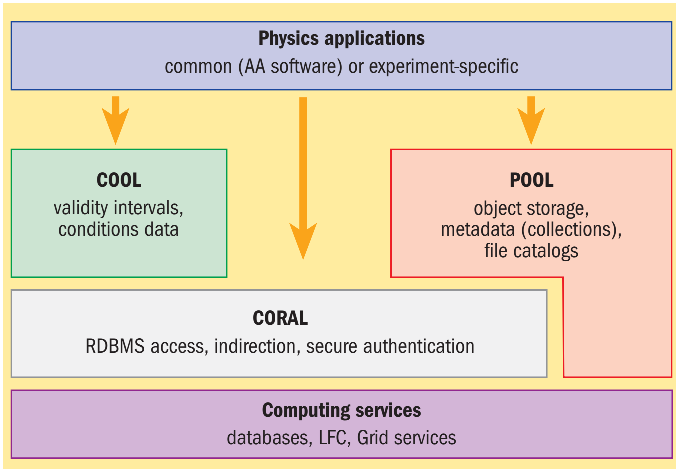


*The LCG persistency framework: layering of software components.*

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 was extended to provide a generic database access layer (CORAL) and a specialized component for storing and looking up conditions data (COOL).

### Accessing databases with CORAL

Database access for all persistency framework components proceeds via the CORAL (COmmon Relational Abstraction Layer) package. CORAL is also being used in production by several LHC experiments directly from their offline and online applications.

CORAL provides a set of C++ interfaces that are independent of the database implementation and therefore enable the same code to be used against a variety of database systems. At the moment Oracle, MySQL, SQLite and FroNTier (a Web-based database caching package) are supported via plug-in libraries that can be loaded at application runtime.

Support for several database implementations is important, not only to minimize the risk of technology binding but also to cover the available database deployment infrastructure across LCG sites. The experiment deployment models foresee the use of Oracle at Tier-0 and Tier-1, and the use of MySQL, SQLite or FroNTier at other Tiers. More details will be available in a forthcoming *CNL* article about the LCG 3D project.

To exploit distributed database resources that are becoming available via the LCG 3D project, CORAL provides secure database authentication and indirection, including retrial/failover across multiple database replicas. CORAL resolves user-defined logical database names into physical connections to database servers that are now available. In the event of network or service problems, CORAL connections will failover to the next available database replica, if necessary at a different site.

CORAL implements several database access optimizations directly (such as row prefetching and the efficient use of server-side cursors) and significantly decreases the user effort to implement others (bind variables and bulk DML operations). These single-client optimizations are complemented by a connection pool that minimizes the number of concurrent server connections from larger applications with several database components and improves access to the database.

### POOL: object storage in databases

The POOL hybrid store functionality is now well integrated in the software frameworks of ATLAS, CMS and LHCb, and has been successfully tested in several large-scale data challenges using object storage in ROOT files.

Building on CORAL, POOL was recently generalized to store arbitrary C++ objects in any of the CORAL-supported database systems. This is particularly useful for calibration and configuration data, which cannot easily be managed in files. With this mechanism objects are decomposed according to their C++ type and stored as rows in relational tables. A set of customizable mapping rules enables the user to steer the automated table generation and to control the mapping of C++ data types to their relational counterpart.

Object-relational storage is now being integrated into the software frameworks of ATLAS and CMS.

### COOL handles conditions databases

The COOL (LCG Conditions Database) package provides a software infrastructure for managing conditions data, focusing on the issue of their time variation and versioning.

The development of COOL began at the end of 2004 to replace several disjointed packages previously developed for MySQL and Oracle. COOL still shares their basic data model for conditions data but is now based on a single code implementation and the same

relational schema for all supported back-ends, thanks to the use of CORAL.

In COOL, measured or calculated detector conditions (such as detector temperatures and alignment parameters) are associated with an interval of validity, the time range to which the stored conditions apply. Groups of similar condition items can be organized in a hierarchical structure similar to a file system. Multiple versions of condition data can be maintained (for example, originating from alternative alignment methods) and can be referred to by tag names (similar to release tags in the CVS code management system).

The COOL software provides a high-level C++ interface to store and retrieve the data according to the most important physics use cases. It takes over most of the physical management of database tables and the creation of indices for fast data access, enabling users to focus on the definition of the experiment conditions and their logical structure rather than on database access optimization. COOL enables users to store data either directly inside the database tables or to maintain references to data stored externally (such as XML or POOL files or other databases), depending on data volume and the experiment deployment model.

COOL is today the baseline conditions data implementation for the ATLAS and LHCb experiments. Although COOL is still being optimized, its performance already matches some of the experiment requirements. Sustained data rates over 20 MB/s and 20 k rows/s have been observed for retrieval from an Oracle RAC cluster database.

**Summary**

With the introduction of the CORAL and COOL packages alongside POOL, the LCG persistency framework is now providing storage functionality for all major physics data types as a consistent set of layered components. Designing these components in close contact with the experiments and the database service providers at CERN and other LCG tier sites will make it possible to successfully deploy both software and services for the start-up of the LHC.

**Further information**

- POOL (the LCG Persistency Framework): http://pool.cern.ch;
- CORAL (COmmon Relational Abstraction Layer): http://pool.cern.ch/coral;
- COOL (LCG Conditions Database): http://cern.ch/cool.

**Radovan Chytracek, Dirk Düllmann, Giacomo Govi, Ioannis Papadopoulos and Andrea Valassi, IT/PSS, CERN**

# The Grid: revolution or evolution?

Why do the planets revolve around the Sun? Has genetic science shaken Darwin's theories to their foundations? Are viruses the champions of evolution? Is progress a form of tradition? On 8 and 9 July, Geneva's Science History Museum held science nights on the theme of "Evolution, Revolution".

CERN took part in the event, anticipating the (r)evolutions from the Large Hadron Collider (LHC). The future accelerator has already led to technological breakthroughs, it promises to deliver more scientific advances, and may even turn our understanding of the infinitesimally small on its head.

The CERN stand put a special emphasis on the computing revolution, from the Web to the Grid. Visitors were able to learn about the Grid in general, about Grid projects at CERN (EGEE, LCG and openlab) and about activities related to distributed computing such as Africa@home. Computer animations explained the technologies, which have spin-offs well beyond the field of particle physics that are of benefit to the whole of society.

Some 30 000 visitors attended the science nights, according to the organizers. Many people flocked to the stands and animations until late on Saturday night and all Sunday afternoon. The CERN stand received much interest and the 15 volunteers from the IT department and the Visits Service were kept busy answering questions. These ranged from the general "How does this work?" to more technical questions like "On what platforms does the Grid run?", and of course "How can I access this?" To make it more hands on, a computer game for children was available in which they could try to save the world by using the Grid to decrypt a message from outer space.

*Adults and children visited the CERN stand at Geneva's science nights.*

So, is the Grid a revolution or an evolution? Historically the Grid can be seen as the latest step in the process of developing distributed computing solutions, an evolution that started many years ago. The switch from mainframe computers to clusters of PCs more than a decade ago is one step in this direction. The significance of the Grid today comes from its ambition to extend distributed computing to a global scale – now possible thanks to high-speed networks – and thus to change the way scientists work together.

In fact, every technological revolution conceals a great deal of evolution. The World Wide Web was itself based on other technologies, such as hypertext, and built on prior efforts to unify electronic information. The same is surely true for the Grid and distributed computing.

**The IT communication team, CERN**

# The DIANE user-scheduler provides quality of service

Today the mainstream use of Grids resembles a large batch system: the goal is to maximize the computational throughput over long periods of time.

This fits many applications, in particular large data productions of the Large Hadron Collider (LHC) experiments, where the production manager puts thousands of jobs into the system and after several days they come out with the result. However, this model does not support other scenarios well. For example, in an interactive analysis the response of the system should be much faster and aligned with the activity of the user; and life-science applications often involve short-deadline jobs, that is many short jobs that must finish within a certain time limit. In general, such quality of service (QoS) characteristics are not present in today's Grid systems.

The scale and complexity of the Grid also has implications. The LHC Computing Grid/Enabling Grids for E-sciencE (LCG/EGEE) is the world's largest Grid system to date, comprising more than 20 000 worker nodes, some 200 computing sites and petabytes of storage. Such an impressive enterprise, which connects heterogeneous computing environments and organizations, comes at a cost: from the end-user's perspective, tracking problems can be time consuming and the system may sometimes be less efficient.

### The DIANE project

User-level scheduling is a light software technique that enables new capabilities to be added and QoS characteristics and reliability to be improved, on top of the existing Grid middleware and infrastructure.

DIANE (DIstributed ANalysis Environment) is such a tool. It is an R&D project for parallel scientific applications in the master–worker model that was started at CERN in 2001. At the beginning the target was to investigate distributed ntuple analysis for particle physics. However, with time DIANE has become an application-independent user-scheduling tool on the Grid (see http://cern.ch/diane). It has been interfaced to a number of applications in high-energy physics, medical physics, life sciences and other fields.

DIANE is a python framework based on a master–worker processing model that is used on top of regular Grid middleware in a transparent way. Worker agents are sent to the Grid as regular Grid jobs. By opening a TCP/IP connection they register to the master agent that runs on the user's desktop computer and is the coordination point for the virtual worker pool. Workers may dynamically join and leave the pool, without disrupting the processing as a whole. The units of computation are many short tasks, which the master allocates to workers directly, bypassing the middleware scheduling layer.

This makes it possible to reduce the total job turn-around time and to react much faster to errors in task execution by reallocating them to other workers. Splitting the processing into many fine-grained tasks improves the load balancing and ensures that the workers are used efficiently. As a result the computing resources may be returned to the Grid faster, because worker agents are automatically terminated when the processing reaches its end.

DIANE's python framework enables existing applications to be integrated quickly, even those as complex as Athena, the analysis framework of the ATLAS experiment. Studies performed by members of the ATLAS collaboration showed that DIANE can be used to integrate local and Grid resources, and even resources from different Grid infrastructures at the same time. The DIANE-based parallel Athena prototype has been shown at EGEE conferences and has been included in the ATLAS Technical Design Report (TDR 2005).

DIANE has also been interfaced to Ganga, a user-friendly Grid interface created in the context of ATLAS and LHCb experiments at CERN. In future, physicists using Ganga will be able to choose the DIANE optimizer, which will be attached transparently to their jobs.

The DIANE scheduler will also be used to operate the statistical regression testing part of the Geant4 release validation procedure. It enables turnaround time to be reduced and provides a more stable and predictable job output rate. This is because the worker agents acquired at the beginning of processing are held inside the pool and are shielded from the instabilities in the Grid brokering. Stable job output rate is an important QoS feature because it enables testing operations on the Grid to be planned with more reliability.

### Practical applications

Earlier this year DIANE was used to perform a sizeable fraction of an *in silico* drug discovery application using the EGEE and other Grid infrastructures. The challenge was to analyse possible drug components against the avian flu virus H5N1.

This activity showed that a user-level scheduler like DIANE can improve the distribution efficiency on the Grid from below 40% to above 80% by optimizing the allocation of the fine-grained computing tasks. Automatic error-recovery mechanisms proved to be efficient in extended periods of continuous work: the part performed with DIANE lasted around 30 days.

In May and June, CERN successfully supported a series of large-scale data-processing activities carried out by the International Telecommunications Union (ITU) as part of the ITU's Regional Radiocommunication Conference. Several sites of the EGEE infrastructure provided a computing Grid of more than 400 PCs to work on each analysis in parallel, and the processing was conducted using the DIANE scheduling layer.

The system completed more than 200 000 very short frequency analysis jobs (clustered in around 40 000 processing tasks) in around one hour, proving that on-demand computing with a short deadline is possible on the Grid. The frequency allocation plan that was optimized with the help of the Grid enabled more than 1000 delegates from 104 countries to adopt the treaty agreement that will replace the analogue broadcasting plans that have existed since 1961 for Europe and since 1989 for Africa.

In the future, closer integration with Ganga will enable access to all of DIANE's capabilities. Ongoing PhD research is aimed at supporting hard QoS requirements with novel techniques such as a floating worker pool, extending scalability above 500 worker agents, and supporting inter-dependent tasks for workflow applications.

**Jakub Moscicki, IT/PSS, CERN**

## LCG launches online bulletin

The *LCG Bulletin* was launched at the beginning of the summer. It is available at https://cern.ch/twiki/bin/view/LCG/LcgBulletins.

The aim of the bulletin is to streamline the distribution of practical information for the LCG community, in the form of short summaries. If you wish to publish information in the bulletin, such as news or future events involving your area of the LCG project, mail this to me with the relevant links.

**Alberto Aimar, IT/LCG, CERN**

# BARC collaborates with LCG

Situated on the shores of the Arabian Sea, on the outskirts of Mumbai in India, lies the Bhabha Atomic Research Centre (BARC). The research centre is an Indian Government facility that employs some 10 000 people, and that has been contributing human resources to the LCG project for the past 21 months. I was one of several members of a CERN contingent from the LCG project who were welcomed to BARC during the CHEP'06 conference in Mumbai.

A remarkable range of projects was presented during the visit. In the area of security, a range of solutions for intrusion detection and incident analysis were discussed, along with a hand-scanning biometric system. This is especially important for BARC, which maintains several isolated networks. Another interesting project was a fast, low-cost 2 million pixel display with 16 separate panels driven by a cluster of machines. This display is used for applications such as satellite imagery, CAD/CAM, medical analysis and tsunami simulation.

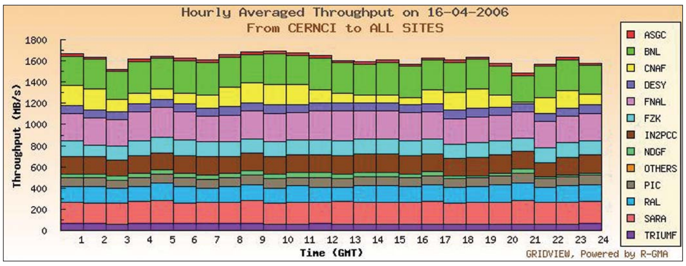


*GridView monitor of data transfer rates from CERN to Tier-1s in April, showing peaks above 1.6 GB/s. The monitor was developed by BARC.*

Areas where BARC has been contributing expertise to the LCG project include the Extremely Large Fabric management system (ELFms), as well as Grid-monitoring technology. As Les Robertson, the LCG project leader, commented during the visit, "the collaboration enables LCG to benefit from BARC's long experience with high-performance computing".

BARC has been instrumental in the development of GridView to support the Grid-level monitoring for the service challenges (http://gridview.cern.ch/GRIDVIEW). BARC has also contributed to LEMON, the lower-level LHC-Era MONitoring component of ELFms, which is being used to monitor 80 metrics across approximately 2500 hosts in 100 clusters. Some 80 automatic recovery actions have now been defined, resulting in a reduction in problem tickets and the associated reduction in human intervention.

CCTracker is a tool used to visualize and plan the physical layout of computer centres such as the LCG Tier-0, and to search for equipment. BARC plans to develop it further to perform high-level service management use cases across sets of nodes and clusters.

Speaking at the CHEP'06 conference, the president of India, Dr A P J Abdul Kalam, referred to the collaboration with CERN in areas such as Grid monitoring and fabric management as a model for how India can contribute to the LHC project.

In addition to the advances in the quality and scalability of fabric-management components, the exchange of ideas, experiences and solutions between BARC and the LCG project has proved beneficial to both parties. BARC is planning to adopt several ELFms components for their operations, including CCTracker and LEMON, and the LCG project is benefiting from the additional skilled resources at this critical time.

**Bill Tomlin, IT/FIO, CERN**

# IN2P3 Computing Centre prepares for the LHC

The IN2P3 Computing Centre (CC-IN2P3), located in Lyon, is the national facility for data storage and processing of the French National Institute of Nuclear Physics and Particle Physics (IN2P3).

Funded by the National Centre for Scientific Research (CNRS) and the Atomic Energy Commission (CEA/DSM/Dapnia), it has provided computing services for more than two decades to several experiments in the fields of nuclear physics, particle physics and astroparticle physics.

Experiments of the Large Electron Positron (LEP) collider, and more recently DZero and Babar, are examples of major users of the centre. Since the early 2000s it has also provided computing services to research institutions in the field of biomedical applications.

CC-IN2P3 has set up and operates a Tier-1 centre to process data from the four Large Hadron Collider (LHC) experiments. It is planned to eventually contribute about 9% of the total worldwide Tier-1 computing capacity for ALICE, 13% for ATLAS, 10% for CMS and 27% for LHCb.

Becoming an LHC Computing Grid (LCG) Tier-1 centre means many changes to the site. A major upgrade of the cooling and power infrastructure of the centre's computer room is scheduled for the second half of 2006. This will enable it to host the data-processing equipment that is essential for the LHC experiments and for other scientific experiments that the centre will continue to support in the coming years. A second machine room in a new building is planned to extend the site's capacity by 2009.

A dedicated optical network circuit that links CC-IN2P3 and CERN at 10 Gbit/s has been in operation since early 2006. It is being used to validate the data-

*The IN2P3 Computing Centre operates a Tier-1 site for the LCG.*

exchange infrastructure for LHC experiments in the context of the LCG project. Data transfer rates of 250 MB/s have been demonstrated between CERN and CC-IN2P3. This circuit, along with the site's links to the national and international networks, are operated by RENATER, the French telecoms network for research and education.

Work has also begun to upgrade the disk- and tape-based data storage and the computing infrastructure of the centre to the levels required for the LCG project. The new Grid components at the site are being integrated progressively into the production-level procedures, with the aim of reaching a high quality of continuous service by the time LHC starts.

CC-IN2P3 has been actively involved in Grid activities since Datagrid, the first European-level Grid project. It also contributes to Enabling Grids for E-sciencE (EGEE), both as a regional operations centre for France and as the developer and operator of the EGEE central daily operations portal.

The LCG, like other large-scale projects, is both a major technological push and a significant human enterprise. In the year of the 20th anniversary since CC-IN2P3 moved from Paris to Lyons, its staff are working hard to meet the challenges presented by the LHC over the coming years.

**Fabio Hernandez, CC-IN2P3**

# USENIX 2006 presents advanced research in computing systems

The USENIX 2006 conference was held in Boston 30 May – 3 June. The annual event, organized by USENIX, the Advanced Computing Systems Association, is a technical conference and a meeting ground for engineers, system administrators, scientists and technicians. This year 800–900 participants attended.

**Extensive programme**

USENIX 2006 offered a wide programme. There were two main tracks: training sessions and technical sessions. The conference included five days of training sessions (full day or half day) and three days of 1.5 hour technical sessions. In parallel with the training and technical sessions there were guru sessions, birds-of-a-feather sessions, invited talks, a poster session and social activities.

The guru sessions were informal meetings where an expert introduced his field and the rest of the time was used for open questions and discussions. Birds-of-a-feather sessions were time slots where anybody could suggest a subject to discuss. These were also informal and many of the subjects were decided during the conference.

The invited talks were 1.5 hour presentations but with an unrelated theme. Some 10 posters were presented at the poster session, which was part of the social activities. The main part of the social programme was the reception held at the Boston Aquarium.

**Training sessions**

More than 30 full- and half-day training sessions were held. The speakers had good teaching skills and were known experts in their fields. Both old and new topics were taught, such as Perl and AJAX, and there was a good range of subjects: from network security, Solaris 10, storage, LDAP, databases, back-up and Linux security to project management and time management. Some of the tutorials are described below.

*Participants learned about the open-source software used to make Pixar Animation Studios' latest film,* Cars.

*Perl*

The Perl tutorial was taught by Tom Christiansen, who has been involved with Perl since day one. He is also the author of several books on Perl, such as *The Perl Cookbook* from O'Reilly. His tutorial started with the basic syntax of Perl, different Perl modules, and Perl as an object-oriented language. He also talked about how to manage exceptions and warnings, and included a good presentation on regular expressions. The tutorial also covered threads and Unicode.

*AJAX*

AJAX is short for Asynchronous JavaScript and XML. It is a Web-development technique for creating interactive Web pages. This is done by exchanging small amounts of data with the server such that the entire Web page does not have to be reloaded each time the user makes a change.

The AJAX tutorial was given by Alex Russel, the project leader for the Dojo Toolkit and primary author of the netWindows DHTML framework. After giving a detailed presentation of JavaScript and the basics of AJAX, he then introduced different JavaScript toolkits with an emphasis on the Dojo Toolkit. The tutorial included an introduction to debugging tools (like Mozilla Firebug), code examples and demos.

*PHP5*

The PHP5 tutorial was taught by David Sklar. He has an extensive knowledge of PHP (Hypertext Preprocessor) and has written several books on the subject, including *PHP Cookbook* and *Learning PHP5*.

In his tutorial he spoke about basic PHP syntax, HTML forms, how to connect to databases with PHP, and how sessions and cookies work. Sklar also offered some tips on how to write secure code, and he introduced some PEAR modules, one being the HTML_QuickForm. He also talked about how to use JavaScript with PHP, and provided code examples throughout.

*Project management*

The project management course was given by Strata Rose Chalup. She has a background in system administration and has had roles ranging from project manager to director. She has written several articles on project management.

In her tutorial she taught the basics of project management, from planning to action. She explained how to group tasks, find dependencies between them and how to do milestoning. Other important subjects were risk management, project tracking, reporting and how to estimate time and work. Chalup also introduced some useful project management tools.

*Time management*

This tutorial was taught by Tom Limoncelli, who wrote *Time Management for System Administrators*. He also cowrote *The Practice of System and Network Administration*.

Limoncelli started by explaining why typical time management books do not work well for system administrators. The main reason is that system administrators have a higher degree of user interruptions and are technical people who have the ability to automate and use tools to help them with their daily work. Limoncelli also talked about how to delegate, handle user requests, make

"to do" lists work and prioritize, and he explained how building routines can save time.

**Technical sessions**

The technical sessions of presented papers covered many interesting areas in computer science, such as virtualization, embedded systems, security and administration.

The programme opened on 1 June with a keynote address by Larry Peterson, the director of PlanetLab. He described the PlanetLab architecture, which comprises a pool of physical machines on which virtual machines (VMs) run. Using VMs, the physical resources are divided between users who can run their networking applications without interrupting each other. The talk set the topic for three sessions held that day (reported below), which were about virtualization and how to fill the semantic gap between physical and virtual machines.

*This year's USENIX conference was held at the Marriott Boston Copley Place hotel, in the heart of Boston's historic Back Bay district.*

*Antfarm*

The Antfarm project by the University of Wisconsin has increased the throughput in performing input/output (I/O) to disk in virtual machine managers (VMMs) by introducing anticipatory I/O scheduling. The presenters explained how anticipatory scheduling is used in Linux to reduce the travel of the disk head and how they could also apply this to VMMs to achieve the same effect with virtual machines.

*Optimizing network virtualization in Xen*

Modern network interface cards (NICs) have features that enable some work involved in processing Ethernet frames, from the CPU, to be offloaded. These include TCP segmentation offload, scatter/gather I/O, and checksum offload. This project, by the Ecole Polytechnique Fédérale de Lausanne, and Rice University, Houston, has carried out work to optimize Xen's virtual NIC. By implementing the offload features in Xen's virtual NICs among other things, this project had speeded up the network throughput in Xen VMs by a factor of four.

*High-performance VMM-bypass I/O in virtual machines*

The virtualized drivers in Xen are a bottleneck in terms of throughput because all I/O requires the intervention of the VMM. A project by IBM in collaboration with the Ohio State University explored an alternative method for performing network I/O in modern network interfaces such as Infiniband. Such technologies often have a facility that enables user applications to bypass the operating system (OS) to use the network interface directly. This is exploited in this project to enable VMs to bypass the VMM.

**Other technical sessions**

Collecting the history or provenance of a file can be useful for many reasons. A paper on provenance-aware storage systems, presented by Harvard University, explored using a storage system with automatic management of provenance. The researchers argued that by tracking syscalls and I/O operations and recording them in a kernel built-in database, the traces can be used efficiently to detect intrusions, debug builds and create scripts that recall compile-time flags and user document editing.

In an enterprise VM system, a lot of redundant VM state is stored because several VMs employ the same OS. Another paper that explored the use of content addressable storage identified redundancy using hashing instead of names in the file system. By eliminating redundancy, network congestion can be reduced significantly.

**Invited talks**

In addition to the paper presentations, there were some interesting invited talks.

Greg Brandeau, the vice-president of technology at Pixar Animation Studios, talked about the studio's new film, *Cars*, and about the computers and software used to make it. Significant parts of the studio's software stack came from the open-source community. The studio wanted to address the open-source community to ask them to write better software. Brandeau also explained some of the routines used to create and render scenes in the animations, including texturing, lighting and physics.

Trevor Blackwell, the chief technology officer from Anybots, showed some videos of the company's robot creations. Using conventional software and computers for controlling movement, Anybots has created a Segway-like unicycle and self-balancing robots that are on wheels and have controllable arms. Blackwell argued that autonomous robots are not feasible at present. Rather, remote high-level control of robots' actions could enable robot control centres to employ a workforce of remote-controlled robot butlers, thus optimizing the controlling human workforce.

**Hege Hansbakk, IT/DES; Håvard Bjerke, IT/DI**

**The deadline for submissions to the next issue of CNL is 6 October**

**E-mail your contributions to cnl.editor@cern.ch**

**If you would like to be informed by e-mail when a new issue of CNL is available, subscribe to the mailing list cern-cnl-info. You can do this from the CERN CNL website at http://cern.ch/cnl**

# CERN will benefit from Google's Book Search and Scholar initiatives

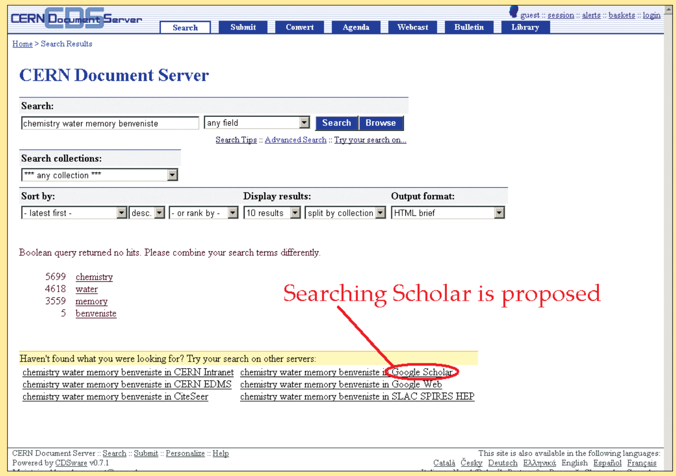


*The CDS Invenio website suggests searching with Google Scholar.*

Jean-Yves Le Meur of CERN IT/UDS, who is responsible for the CERN Document Server (CDS), talks to *CNL* about Google's recent projects, Google Book Search and Google Scholar. In the following interview Le Meur explains how these initiatives may increase the dissemination of CERN documents to a wider audience, and how they may improve searches within CDS of CERN's own book collection.

**Can you explain the differences between Google, Google Scholar and Google Book Search?**
Everybody knows Google Web, the search engine for Web pages worldwide.

Google Web follows all links and it indexes all public information but it cannot systematically reach the "deep Web", which mostly comprises databases of documents.

With Google Scholar (http://scholar.google.com), which is still in Beta version, Google has begun to address the issue of retrieving the content of the deep Web. The goal is to provide a simple way to search for scholarly literature in all disciplines, with results sorted by the most cited documents first.

The databases of documents indexed by Google Scholar are composed of peer-reviewed papers, theses, books, abstracts and articles, from academic publishers, professional societies, preprint repositories, universities and other scholarly organizations. Google Scholar provides a single entry point that enables people to search for public professional information across many disciplines and sources – a kind of "network of excellence" search engine.

Google Book Search (http://books.google.com), which is also still in Beta version, aims to establish a search engine for the huge corpus of traditional paper books. Google is collaborating with publishers and libraries to digitize collections. This will enable everyone to query the full text of books to find ones of interest and learn where to buy or borrow them. For old books in the public domain, the full text may also become freely available via open access for reading.

**What is the relationship between a document management system like CDS and Google Scholar?**
CDS is a digital library that aims to collect all high-energy physics (HEP) and HEP-related literature. Of course, it is also used as the CERN institutional repository in charge of archiving and disseminating documents produced by the organization. With the birth of Google Scholar and the quality of the citation ranking it provides, we at CDS are keen to ensure that HEP and CERN works are well covered by this new service. People searching in Google Scholar should never miss documents from CERN.

In collaboration with the CERN library and with Google, we are setting up procedures to include in our CDS Invenio software (at http://cdsware.cern.ch/invenio) the mechanisms needed to guarantee that in the future Google Scholar indexes and ranks all new and updated documents from CDS. As Google Scholar has based its harvesting policy on the National Library of Medicine's (PubMed's) Journal Publishing Document Type Definition, we are now in the process of converting CDS information into this format. Our own search engine already proposes to users who do not find what they are looking for that they search on Google Scholar, where articles from other disciplines can also be retrieved.

**Do you think Google Book Search will modify the way physicists gain access to information?**
Yes. The Google Book Search initiative to create a worldwide digital library has triggered multiple reactions in the EU, where some fears were expressed that the selection of books would rely on a private US company. This suddenly very active environment for the creation of digital libraries means that in the medium or long term, worldwide access to all books, including physics books, will be much easier through Google Book Search, and probably also through other similar services.

We and the CERN library have been talking to Google Book Search representatives to find out how best CERN can take advantage of this initiative. For example, each time a book available in the CERN library is scanned and indexed by Google Book Search (which deals directly with physics publishers), CDS would like to receive its ISBN number so that we can point to Google Book Search in exactly the same way we already do with Amazon.

Another example of an add-on service would be to use the full-text indexes of books belonging to the CERN collection, created by Google through optical character recognition, to provide a better book search facility within CDS on top of the library holdings.

**Conclusion**
Google Scholar and Book Search are ambitious projects that are still at the prototype stage. We have already started to actively collaborate with them because we believe that they can improve the search experience of scientists looking for academic literature. Creating gateways between CDS and these new information retrieval services is therefore a natural way to better serve the HEP community.

# Migration of the printing service is under way

The print server infrastructure at CERN is now in a phase of migration. The migration is planned to be transparent for users and there should be no interruption to service.

During the transition period some of the printers will be hosted by the new print servers while others will be connected to the old ones. Therefore it is important to find out which server a printer is configured to before troubleshooting or looking for documentation. A different procedure has been introduced for Windows users wishing to add new printers, depending on the server a printer is connected to.

You can check the status of your printer at: http://cern.ch/WinServices/Services/Printers. Type the name of your printer in the upper text box and click the Search button (see figure).

You can also consult the list of printers that have already been migrated at: http://cern.ch/WinServices/Services/Printers/PrinterList.aspx.

## Documentation

If you have a Windows computer and want to use a printer that is still connected to the old service, please follow the usual documentation of the CERN Printer Wizard at http://cern.ch/WinServices/Help/?fdid=7.

For other cases the following documentation is available:

- NICE XP/2000 – http://cern.ch/WinServices/Help/?kbid=070103
- SLC3 – http://cern.ch/linux/scientific3/docs/printing.shtml
- SLC4 – http://cern.ch/linux/scientific4/docs/printing.shtml
- Mac OS X – http://cern.ch/it-dep/gencomputing/mac-support/PrinterSelect.htm
- Visitors (Windows) – http://cern.ch/WinServices/Help/?kbid=070107#visitor
- Visitors (Linux) – http://cern.ch/linux/documentation/printing.shtml

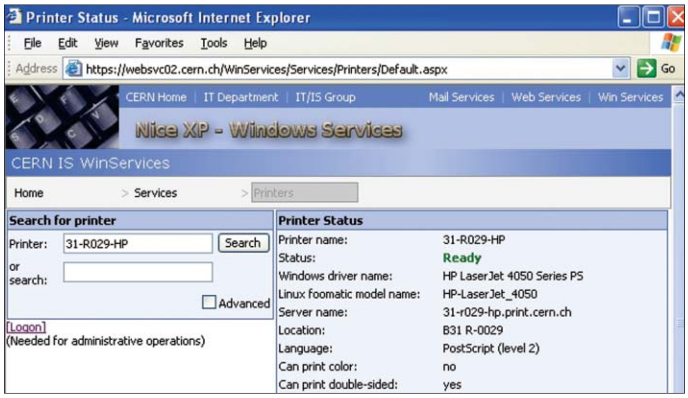


*Users can check the status of their printer at Windows Services.*

## Operating systems

The following information about the eventual impact of the migration depends on the operating system you use.

*NICE XP/2000*
The printer migration on the NICE environment will be completely transparent. It will be carried out automatically on each computer by moving already configured printers from the old to the new service according to the printer status at a given time. The migration job will also track changes that users might have made manually in the meantime, also in a completely transparent way.

*Scientific Linux CERN 3 and 4*
The migration has already been completed for Linux computers. All printer connections were reconfigured to use client-side processing of print jobs. This reconfiguration was done automatically during the system update of SLC3/4 on the night of 31 July. If necessary you can reconfigure your system again at any time by running: /usr/bin/cern-config-printers –u. Please note that support for "Xprint" has been dropped on the new system. The "lpr" command should be used instead.

*Mac OS X*
Migration is also transparent. Nothing will change for Mac OS X computers but please ensure that the printer connections are configured in the way described under the following link: http://cern.ch/itdiv/gencomputing/mac-support/PrinterSelect.htm.

## Further information

For more details please consult http://cern.ch/WinServices/Help/?kbid=070001. We also recommend the FAQ section at http://cern.ch/WinServices/Help/?kbid=070020.

A more detailed article about the new printing infrastructure will be published in the November issue of *CNL*. If you have any questions please contact printer.support@cern.ch.

**Rafal Otto, IT/IS**

# Request from the helpdesk: please don't thank us

This may seem a little unusual but recently we have started to explicitly ask users not to say "thank you" when they are pleased with a solution provided by the computing helpdesk. This message is now included in the mail that we send when providing a possible solution (see figure).

The explanation for this request is that we want to automate the reopening of tickets from users who inform us that the solution is not totally satisfactory. The system will automatically reopen the same ticket number and a follow-up will be made by the support staff. However, it is not possible for the system to distinguish between a request for more help and a "thank you".

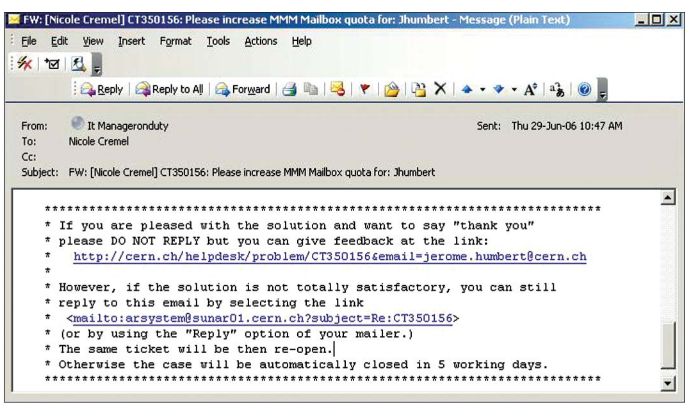


*Extract of the e-mail sent to a user when a solution has been provided.*

Users can still express their satisfaction, however. Instead of sending a direct reply they can access the Web page associated with the ticket number (beginning: "http://cern.ch/helpdesk/problem/CT"...) that is included in the mail they receive. On this page users can see the details of the problem and its treatment, and at the top of the page they are offered the chance to rate the way in which the problem has been handled. Five levels are proposed (very poor, poor, average, satisfactory, very satisfactory) and users can simply select a level.

So please remember not to thank the helpdesk, and instead take a few minutes to open the link and give some quick feedback. Users who wish to provide more comments but do not expect the ticket to be reopened can e-mail the manager on duty at mod@cern.ch.

# Dinosaurs, browsers and monkeys

At the beginning of August the Mozilla application suite, which includes a Web browser and e-mail client, was replaced by a product called SeaMonkey. While the operation was supposed to be largely transparent (the look and feel are similar, and both applications are launched via /usr/bin/mozilla), several users have noticed the change. This article will shed some light on the history behind this.

**In the beginning**

Mozilla was initially the internal codename for a Web browser from Netscape. The name reputedly comes from "Mosaic killer" – Mosaic was an even older browser. Netscape's browser acquired several other utilities, like a mail client, an HTML editor and a chat client. The company decided to hand the source code for this browser application suite for further development to the Mozilla Foundation, a not-for-profit organization, and "Mozilla" became the product's official name.

Some time later new stand-alone versions of the integrated clients received much more attention from developers (the results are the Firefox Web browser and Thunderbird e-mail client), and the integrated Mozilla product did not benefit from newer developments.

The Mozilla Foundation decided to stop developing Mozilla itself, but provided source code and infrastructure help to a new project that aimed to keep the old Mozilla alive and secure. To avoid confusion between the Mozilla Foundation and the integrated browser plus e-mail application, the latter was renamed SeaMonkey (this also was originally an internal codename at Netscape). More documentation can be found on the SeaMonkey project pages.

**SeaMonkey arrives**

Recently a spate of new vulnerabilities in the old Mozilla, that were promptly fixed in SeaMonkey, led Red Hat – a major Linux distributor – to replace the Mozilla package with the new SeaMonkey. Scientific Linux followed suit. Care was taken to keep the two applications compatible wherever possible, so that configurations, shortcuts and browser plug-ins would still work.

Such project renames are quite common in the open-source community; often a clashing trademark is only discovered some time after a project has been launched. For example, the Firefox browser was initially called Phoenix but a PC BIOS producer of that name exists. It was then briefly called Firebird, but changed to Firefox to avoid confusion with the Firebird free database software.

Projects also have their own release cycles, so sometimes major updates or even product replacements fall into the lifetime of a distribution and trigger major updates. For example, Firefox 1.0 was recently updated to 1.5 in both stable versions of Scientific Linux CERN, since 1.0 was no longer maintained.

Such hiccups are unavoidable since Linux distributions consist of the results of several hundred (or several thousand) individual development products, each with its own release schedule.

**Jan Iven, IT/FIO, for the CERN Linux team**

# Computing Colloquia focus on mobile communications

CERN Computing Colloquia present future trends in computing and information technology that are of broad interest to the physics and computing community at CERN. The colloquia complement the more technical IT seminars, also organized by the IT department.

The autumn/winter series of colloquia are on the theme of mobile communications. There are now over two billion mobile phones on the planet, and a range of related handheld devices are changing the way people work and live. The Computing Colloquia will host a series of distinguished speakers from leading companies in the mobile communications sector who will discuss the future of this fast-moving technology.

*Distinguished speakers will talk about future trends in computing.*

**Beyond SMS: what are the "killer applications" for mobile communication devices in the next five years and who will develop them?**

Prof. Jan Bosch, Head of Software and Application Technologies, Vice-President, Nokia Research Centre, Helsinki, Finland.
Wednesday 13 September, 14:00–15:00, Council Chamber.

**Fundamental science and technological progress: links between theoretical physics and mobile communication technology**

Mike Lazaridis, founder and co-CEO, Research in Motion; Chancellor of the University of Waterloo, Canada.
Tuesday 24 October, 14:00–15:00, Main Auditorium.

**Semiconductor technology for a wireless world: tackling the dual challenges of power consumption and performance improvement**

Sir Robin Saxby, Chairman of ARM; President of the Institution of Engineering and Technology.
Friday 26 January 2007, 11:00–12:00, Council Chamber.

The colloquia are free and open to all. To participate from outside CERN, please contact the IT Communications Team (mailto: Francois.Grey@cern.ch). Details of how to propose a colloquium are available at http://computing-colloquia.web.cern.ch/computing-colloquia/howto.html.

# Calendar

**September**

19–21 **European Grid Technology Days 2006**
Brussels, Belgium
http://cordis.europa.eu/ist/grids/event-190906.htm

21–23 **2nd Austrian Grid Symposium and 6th Austrian-Hungarian Workshop on Distributed and Parallel Systems**
Innsbruck, Austria
www.austriangrid.at/symposium

22–24 **6th WSEAS International Conference on Simulation, Modelling and Optimization (SMO '06)**
Lisbon, Portugal
http://www.worldses.org/conferences/2006/lisbon/smo

25–27 **High-Gradient RF 2006**
CERN, Geneva, Switzerland
http://cern.ch/HG2006

25–28 **Cluster 2006**
Barcelona, Spain
http://cluster2006.org

25–29 **EGEE Conference**
Geneva, Switzerland
http://www.eu-egee.org/EGEE06

28–29 **Grid 2006**
Barcelona, Spain
http://people.ac.upc.edu/rosab/grid2006

**October**

9–13 **HEPiX Autumn 2006**
Jefferson Lab, Virginia, US
http://conferences.jlab.org/HEPiX

12–15 **8th Real-Time Linux Workshop**
Lanzhou, China
www.realtimelinuxfoundation.org/events/rtlws-2006/ws.html

21–23 **GCC 2006**
Changsha, China
www.vce.org.cn/gcc2006

25–27 **eChallenges e-2006**
Barcelona, Spain
www.echallenges.org

**November**

11–17 **SC06**
Tampa, USA
http://sc06.supercomputing.org