

CERN COMPUTER NEWSLETTER

Volume 42, issue 1 January–March 2007

Contents

Editorial

CS group connects physicists to LHC 1
Changes to the Computer Newsletter in 2007 2

Announcements & news

Users may need to request deletion of mailbox 3
Support ends for old Microsoft programs 3
Computing articles featured in this month's *CERN Courier* 3

In memoriam

David Owen Williams 1944–2006 4

Physics computing

SLC4 becomes default Linux version on PCs 5
Physicists push limit of network data transfer 5
Chip boosts physics computing power 5

Grid news

WISDOM continues the battle against malaria 6
IT launches Gridcast at Supercomputing'06 6
The Grid improves efficiency of Geant4 production cycles 7
CERN Glite Grid community uses virtualization 8

Technical brief

CERN CA provides security 9

Conference & event report

HENP community shares technology at HEPiX meeting 12
CSC2006 sets three new records 13

Viewpoint

TWiki facilitates collaborative information handling at CERN 14

Information corner

How to help keep you PC secure 15
CERN Computing Colloquia Spring/Summer series kicks off with open source software talk 15
openlab invites applications for its latest student programme 16
Recent changes to IT services 16
NICE streamlines printing infrastructure 16
Recent acquisitions at the IT bookshop 16
Solaris service stops 16
Calendar 16

Editors Nicole Crémel and Hannelore Hämmerle, CERN IT Department, 1211 Geneva 23, Switzerland. E-mail: cnl.editor@cern.ch. Fax: +41 (22) 766 8500. Web: cerncourier.com/articles/cnl.

Advisory board Wolfgang von Rüden (head of IT Department), François Grey (IT Communication team leader), Christine Sutton (*CERN Courier* editor), Tim Smith (group leader, User and Document Services).

Produced for CERN by IOP 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

©2007 CERN

The contents of this newsletter do not necessarily represent the views of CERN management.

IOP Publishing



CS group connects physicists to the LHC



Some of the networking equipment at the heart of the LCG. (Courtesy Wim Klerkx.)

In 2007, *CNL* will continue to talk to group leaders in the IT department about their groups' current activities and future plans. Throughout the year *CNL* will focus on groups that are providing key support to the LHC Computing Grid (LCG), and the challenges they are facing as physicists prepare to analyse real data.

We start with an interview of communication systems (CS) group leader, David Foster. The four major services that the CS group is responsible for are the internal networking infrastructure on the campus, the networking for the experiments and accelerators, the external networking and telephony.

How is the campus network at CERN evolving in preparation for the LHC?

2006 was the first year of a three-year rejuvenation plan for the campus network. There are 150 star points on the site and nearly 800 subnets that need to be updated, as well as over 14 000 cables to be replaced.

We have completed about 30% of this work during the last year. Part of the challenge is the removal of old equipment that has been there for over 10 years.

Most of this work is handled by external contractors, but the CS group provides the overall coordination. This means planning, scheduling and managing the contracts, which is a major activity for the group. We must also monitor and maintain the network once it is in place. A great deal of effort in the group goes into automating the maintenance of the network, including developing software tools that help with this. In addition, we are in the process of improving the level of security for access to the network.

What is the role of CS in supporting the experiments?

We are providing and running a large fraction of the networking infrastructure for the experiments. This includes the cabling for the computing farms in the pits and, with the exception of LHCb, the networking for the control and data acquisition systems. Cabling is completed on ATLAS and LHCb, and over 70% complete on the other two LHC experiments. The necessary routers and switches are being installed.

There is also the LCG computing farm network in the computer centre, which evolved a lot last year with the installation

of eight 10 Gbit/s routers and nearly 200 switches, as well as more than 50 10 Gbit/s router backbone ports and a similar number of 10 Gbit/s router user ports.

What challenges lie ahead as far as high speed connectivity to the external network is concerned?

The 10 Gbit/s connectivity from CERN to the Tier-1 centres is largely complete. In fact, during all the service challenges, network capacity between Tier-0 and Tier-1 was never the limitation.

A significant event in 2006 was that a point of presence (PoP) for the GEANT network was established at CERN. GEANT provides interconnectivity between national networks – including multiple 10 Gbit/s dark fibre connections – and CERN is now the PoP for Switzerland. For internet protocol (IP) connectivity the lab still connects via SWITCH, the Swiss National Network. But for the end-to-end lightpaths which are critical for Tier-1 connectivity, having the PoP at CERN means we can connect directly to GEANT for this service.

Looking further into the future, there is debate in the industry about whether to switch to 40 Gbit/s or 100 Gbit/s. Either way, we will need to start gaining experience with the next generation of equipment over the next few years, so we will be prepared when it becomes competitive. Our ability to test 10 Gbit/s technology well before it became a “commodity” technology played an important role in our being able to deploy it effectively once the price was right.

What impact do you anticipate the LHC start-up will have on the way the physics community uses the wireless infrastructure at CERN?

We have to be prepared for a large influx of users to the site. These people will want to access data from restaurants, the hostel – anywhere. So one thing we anticipate is a large increase in demand for wireless connectivity on campus.

We installed 144 base stations last year for a total on-site of 386. Longer term



David Foster, the group leader of IT/CS.

we're also anticipating a big increase in bandwidth requirements because over the next few years physicists will have so much capacity on their laptops that they will be able to do analysis locally. This means users will want to download perhaps up to 1 TB of data at a time.

Currently we have several implementations of the 802.11 standard for wireless communication on campus. The typical problem is that once you have a lot of ad hoc solutions for wireless around, you get interference. We've reached the point now where this needs to be a managed infrastructure. So in 2007 we'll be looking at technologies and tools that can help us with this task. We want to plan for a homogeneously managed wireless infrastructure throughout the lab.

How is IP telephony affecting CERN?

Although CERN has very competitive rates for traditional telephony the industry

trend is for the telephony infrastructure to be increasingly based on IP. Also, one of the strong user requirements we get is to support incoming IP telephony. So we are looking into technology for accepting and routing IP calls on the site. IP telephony tools are still emerging – the marketplace is not very stable – so this is quite a challenge. Still, we recognize that CERN cannot remain static and must invest in adapting to the external environment.

That said, the lab's traditional telephony services continue to be highly appreciated. Around 270 000 calls were handled by the switchboard last year (one every 90s) as well as 5800 audio conferences. The mobile phone infrastructure also continues to evolve. The Telecom lab handled some 600 new users last year for a total of 3800. In the experiments and accelerator tunnels installation of “leaky feeders” – which provide GSM and VHF communications – is 90% complete. Further down the line, we are preparing for the move to the new UMTS standard, which will have an ultimate capacity of up to 100 Mbit/s compared with under 300 kbit/s for GSM.

What is the future of the CERN Internet eXchange Point (CIXP)?

CIXP, which is operated by the CS group, provides peering between internet service providers (ISPs) and telecoms operators on a commercial basis. Because CERN was an early mover in this business, there is a critical mass of customers. Two new ones were added last year for a total of 32.

Looking to the future of internet exchange points, optical light path exchange is on the cards. This means not just switching IP packets but also light paths – in other words, layer-1 circuits. Operators might want to provide connectivity between places on a temporary or semi-permanent basis or even perhaps on a short-term basis, say to download a large data set. There are a lot of practical problems with this vision, but it is certainly an emerging direction.

Changes to the Computer Newsletter in 2007

For over two years the *CNL* editorial team has collaborated with IOP Publishing (IOP) on the production of *CNL*, as well as on the bi-monthly “Computing News” section in *CERN Courier*. Our readers have expressed satisfaction with the resulting more professional look of *CNL*. But there is always room for improvement.

In reaction to the most frequent feedback we received from readers in 2006 we have moved to a three-column layout, which should improve readability of the text. Also during 2007 IOP is planning to overhaul the website, which should make it easier for users to access and download articles. We

will inform our readers when these changes have been implemented.

Finally, in an effort to disseminate information about Grid developments more effectively around the world, the Enabling Grids for E-science (EGEE) project has joined forces with the US Open Science Grid to produce an e-newsletter, *International Science Grid this Week (iSGTW)*. This is the successor to the *Science Grid This Week* newsletter that has been produced at Fermilab since 2005.

As part of this collaboration *CNL* will feature a Grid news section. This replaces the LCG News section and broadens the Grid news

coverage in *CNL* to include articles related to EGEE and other Grid projects in which CERN is involved. The section will reproduce articles from *iSGTW* that are deemed to be of particular interest to CERN users, as well as continuing to carry articles specific to LHC Computing Grid activities that have been submitted to the *CNL* editors. We feel this change will better reflect the Grid activities in which CERN is involved and that are of interest to CERN users.

CNL readers keen to receive the entire *iSGTW*, which is distributed weekly as an embedded HTML e-mail, are invited to subscribe at www.isgtw.org.

The CNL editors

Users may need to request deletion of mailbox

As a follow-up to the article “E-mail addresses at CERN” published in the last CNL (see www.cerncourier.com/articles/cnl/3/11/8/1) we would like to recommend external users to have only one mailbox. This mailbox could be located either at CERN with the physical e-mail address (PEM) defined as `login@mail.cern.ch`, or at their home institute with the PEM set to this external location.

External users who still have two mailboxes (one at CERN centrally hosted on MAILSERV and another one at their home institute with the PEM defined to the external location) may discover that their NICE account gets blocked regularly without any notification. They notice this, for instance, when they try to connect to CERN via the CERN terminal server. Alternatively, they may notice this when they cannot get access to a Web page that requests NICE credentials.

This occurs for external users with a

Type	Service	Name	Uid	Group/Gid	Last access	Status	Blocking reason	Quota	Shell
Account	MAILSERV	icremel	558	US: 2227	2006-12-10	Active			
Account	AFS	icremel	558	US: 2227	2006-12-08	Active			
Account	AIS	icremel	559	US: 2227	2006-12-05	Active			
Account	NICE	icremel	558	US: 2227	2006-12-10	Active			

Data displayed by CRA when you select “My Details”, including the definition of your e-mail address identified as either a CERN MAILSERV account or other if it is external. Remove by selecting the red cross.

PEM defined outside CERN who still have a mail account (and login) at CERN. Since the introduction of the Computing Resources Administration (CRA) system in 2006, if a CERN mailbox is never accessed by a user

(for example all e-mails are forwarded automatically) this will cause blockage of the NICE account.

To avoid this problem we strongly recommend all users who are in this situation to request a deletion of their CERN mailbox. This can be done in one of the following ways:

- Access CRA at <http://cra.cern.ch/> using AIS login and password.
- Contact the computing group administrator.
- Send an e-mail to user.registration@cern.ch.

Deletion of the CERN mailbox does not affect the definition of the CERN address, `Firstname.Lastname@cern.ch`. This is automatically created as soon as a user registers at CERN and is known by the human resources database. If the user’s PEM is external the CERN address simply points to the user’s external mailbox.

Nicole Crémel, IT/UDS

Support ends for old Microsoft programs

With the approval of the Desktop forum, support for Windows 2000, Office 2000 and Office XP at CERN stopped at the end of December 2006. As a consequence of this security patches will no longer be provided, nor will services for reinstallation be available for these programs. From January 2007 only Windows XP and Office 2003 will be supported.

Users still running Windows 2000 have been notified by e-mail and have been advised to install Windows XP/Office 2003 (as described on <http://cern.ch/Win/Help/?kbid=100001>). Users of Office XP who have not yet upgraded to Office 2003 will be upgraded automatically to Office 2003 as of January 2007.

To ensure compatibility between the new Microsoft Office 2007 and previous releases (2000, XP and 2003) the “MS Compatibility Pack for Office 2007” is available for installation on all CERN standard PCs. This pack allows users to save documents in the new format from older versions of Office. More information is available at <http://cern.ch/win/Help/?kbid=030171>. This pack is planned to be installed on all CERN PCs at the beginning of 2007.

Users who absolutely need to keep Windows 2000 or other obsolete software are responsible for maintaining their system and making sure it follows the CERN security rules – otherwise the machine will be disconnected from the network.

The NICE Team

Computing articles featured in this month’s CERN Courier

The articles listed below appear in the January–February 2007 issue of *CERN Courier*. Full-text articles and the rest of the issue’s contents are available at www.cerncourier.com.

Computing News

- **Record numbers attend EGEE’06** More than 700 participants discussed the status and future of Grid technologies.
- **SC’06 unveils Grid computing ‘mappa mundi’** One of the first attempts to show the whole scientific Grid world.
- **Physicists push the limits of network data transfer** Record-breaking achievements at SC’06.
- **Science Grid This Week goes global** Building on the success of a weekly e-newsletter.
- **NSF supports collaboration to reinforce Pacific Rim Grid** Accelerating the use and interconnection of Grids around the region.
- **Swiss join forces to capitalize on Grids** Providing a central hub for collaboration and knowledge dissemination.
- **High-speed communication link connects Milan and Mumbai** DANTE and ERNET enable scientific collaboration.
- **L’IN2P3 fête ses 20 ans à Lyon** IN2P3/CNRS a célébré l’anniversaire de son implantation à Lyon-Villeurbanne.

Calendar of events

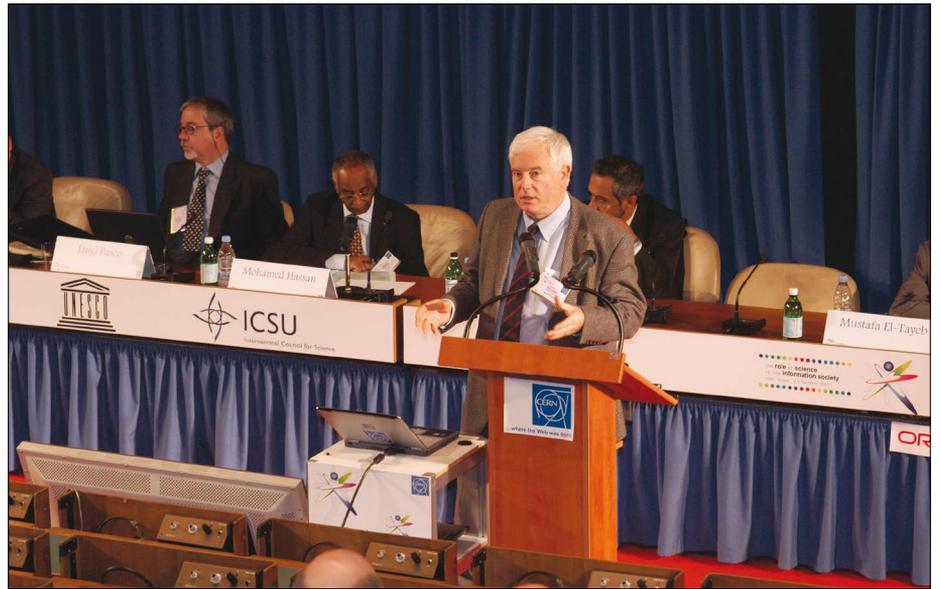


Product information

Feature articles

- **Workshop engages PCs in accelerator controls** Matthew Bickley reports from a meeting at Jefferson Lab that took place on 24–26 October 2006 and brought together accelerator-controls experts from all over the world.

David Owen Williams 1944–2006



Left: David Williams was a senior scientist and coordination officer in relations with the European Union. Right: David at the Role of Science in the Information Society (RSIS) conference in December 2003. The conference was organized to help address the inequalities in scientific activities between advanced and developing countries, a challenge issued by United Nations secretary-general Kofi Annan.

Many people, not only at CERN but also throughout the world, were saddened to learn that their friend and colleague David Williams had died from cancer on 24 October.

David came to CERN from the University of Cambridge in 1966, with a degree in physics and a Masters in computer science. At the time he was one of the few people in the laboratory with formal computer training, as opposed to the physicists who learned on the job. These were very early, if not pioneering, days for computer-science courses, and Cambridge was among the front-runners in the subject.

Joining what was called the Documents and Data (DD) Division, David worked first on software for analysing bubble-chamber photographs. He subsequently led the group that supported experiments with hybrids of bubble chambers and electronic detectors, and then the group supporting online computing in experiments. He thus witnessed the enormous changes that took place in particle physics as the era of bubble chambers came to an end and the powerful, compact mini-computers revolutionized the data collection.

Armed with this experience, David moved on to become deputy leader of DD in 1985 and then leader of the re-named Computing and Networks Division in 1989, the year that the Large Electron–Positron Collider started up. He held this position until 1996, when the member states approved the Large Hadron Collider and its experiments.

This was the period of transition from

central mainframes and supercomputers to a completely distributed computing environment for the acquisition, processing management and analysis of the hundreds of terabytes of data produced by CERN experiments and to a reasonably coherent desktop environment for the more than 10 000 staff and researchers worldwide who use CERN’s computing facilities.

It was also the period in the division when Tim Berners-Lee and collaborators, with David’s strong support, created the World Wide Web. As division leader, David had the foresight to encourage studies that ultimately led to the programming paradigm shift from Fortran to C++.

The period 1996–7 was a time of change for David. He actively fostered the development of the Internet in Europe, not just as a tool for scientific research but also as a motor for Europe’s overall economic development.

This naturally led to his holding a number of positions at national and European level. He was president of the Trans-European Research and Education Networking Association (TERENA) from 1999 to 2003, making a pioneering contribution through the Study into European Research and Education Networking as Targeted by e-Europe (SERENATE) study and developing a strategic vision of the future of research and education networking.

He was also a member of various UK committees dealing with e-science and was chair of the e-Science Advisory Board of the UK’s Council for the Central Laboratory of the Research Councils.

With this broad view, coupled with his energy and managerial talents, it was natural that David moved on at CERN to be responsible for the overall coordination of CERN’s relations with the European Union, a post he held until the end. He played a highly significant role in formulating and obtaining European Commission support for the European part of the LHC worldwide computing environment, which is now almost ready for the LHC start-up in 2007. He was able to persuade the Commission to admit CERN to the Framework Programme (FP), and helped to formulate and organize the FP5 and FP6 projects European DataGrid and Enabling Grids for E-science in Europe.

For those who knew him, David was always a private person, unwilling to talk about himself or worry others with his problems. However, the well merited award in 2005 of an honorary professorship at Edinburgh University was a source of enormous personal satisfaction to him. Many people will remember him as being generous and supportive, always ready to spring to the defence of someone who he considered to be wrongly criticized. He also had the habit, when attending scientific talks, to ask the questions that others were afraid to ask. Although David had usually understood perfectly, this gave the speaker the opportunity to clarify points that had previously been somewhat obscure.

Much sympathy goes out in this difficult time to David’s wife Lidy, to his children, Mark and Marietta, and their families.

Colleagues and friends

SLC4 becomes default Linux version on PCs

Scientific Linux CERN 4 (SLC4) has been the default CERN-supported Linux version for desktop PCs since March 2006 (as announced in *CNL* June–August 2006, available at <http://cerncourier.com/articles/cnl/3/6/16/1>). While we await availability of the gLite middleware, the migration of LXPLUS (with the plan to change the default “lplus” alias to point to SLC4 machines) and the migration of the capacity of LXBATCH to SLC4 has been delayed. However, the migration must take place early in 2007 to ensure a stable environment well in advance of the many computing system tests planned in the run-up to initial LHC data taking.

Since the certification of SLC4 in May 2006, LXPLUS and LXBATCH capacities have been available for both possible

architectures, 32 bit (i386; LXPLUS alias “lx32slc4”) and 64 bit (x86_64; alias “lx64slc4”). To maintain flexibility for the future, production SLC4 capacity will be provided in 64 bit mode (x86_64) using the 32 bit compatibility libraries to allow existing applications to compile and run on this OS without change. This will help to ease both system management and a possible migration of user applications to 64 bit mode.

The LXPLUS and LXBATCH capacities in x86_64 mode have been increased during the last few months. The new modes that will be added to LXBATCH at the beginning of 2007 (some 500 – almost doubling the current capacity) will be available only under SLC4.

A key date in the migration schedule

is the date for the move of the LXPLUS alias from today’s SLC3 service to the new SLC4 cluster running in 64 bit mode. As I write, this is planned for 29 January 2007. However, this change will only be made if the full environment, including the grid middleware, has been properly tested on both LXPLUS and LXBATCH. We will, of course, make sure that any changes to the schedule are published widely. To continue using an SLC3 service after the switch of the LXPLUS alias, please use the alias “lxslc3”.

Finally, I would like to remind all Linux users that support for SLC3 will be stopped by the end of next year. Support for the grid middleware for SLC3 will end in the summer of 2007.

Thorsten Kleinwort, IT/FIO

Physicists push limit of network data transfer

An international team led by the California Institute of Technology (Caltech) joined forces to set new records for sustained data transfer between storage systems during the Supercomputing 2006 (SC’06) Bandwidth Challenge in Tampa, Florida. The team involved high-energy physicists, computer scientists and network engineers from Caltech, CERN, the University of Michigan, the University of Florida and Vanderbilt, as well as participants from Brazil (Rio de Janeiro State University, UER), and the State Universities of São Paulo, USP and UNESP) and Korea

(Kyungpook National University, KISTI).

The demonstration was called “High Speed Data Gathering, Distribution and Analysis for Physics Discoveries at the Large Hadron Collider”. It achieved a peak throughput of 17.77 Gbit/s between clusters of servers at the show floor and at Caltech.

Following the rules set for the SC’06 Bandwidth Challenge, the team used a single 10 Gbit/s link provided by National Lambda Rail that carried data in both directions. Sustained throughput during the night before the bandwidth challenge exceeded 16 Gbit/s using just 10 pairs of small servers sending data at 9 Gbit/s from Tampa to Caltech, and eight pairs of servers sending 7 Gbit/s of data in the reverse direction.

One of the key advances in this demonstration was the use of Fast Data

Transport (FDT), a Java application developed by Caltech that enables stable disk reads and writes coordinated with smooth data flow across the long-range network. FDT streams a large set of files across an open TCP socket so that a large data set of thousands of files, as is typical in high-energy physics applications, can be sent or received at full speed, without the network transfer restarting between files.

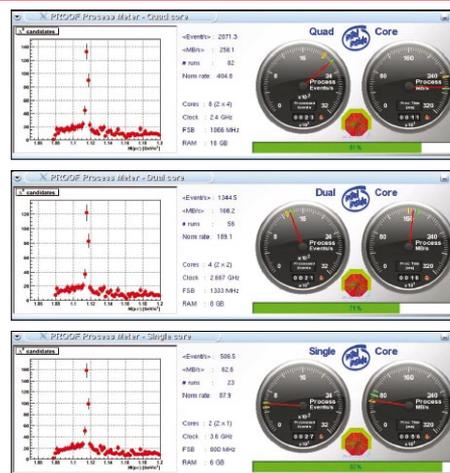
By combining FDT with FAST TCP and an optimized Linux kernel known as the UltraLight kernel, the team reached unprecedented throughput levels, which were limited only by the speeds of the disks. This corresponded to reading at 9 GB/s, or writing at 5 GB/s to a rack of 40 low-cost servers.

François Grey, IT/DI

Chip boosts physics computing power

During an Intel media event at CERN on 14 November 2006 to launch Intel’s new quadcore processor, CERN openlab researchers presented the first tests of this chip for physics computing applications.

The strategy of putting multiple processor cores on the same chip, which has been pursued by Intel and other chip-makers in the last few years, represents an important departure from the more traditional improvements in the speed of such chips. The steady rise of the clock-frequency of processors, a rise which propelled the industry for decades, was resulting in major design challenges. A widely quoted comparison was that the density of heat produced on a multi-GHz chip was approaching that produced in a nuclear reactor, leading to engineering problems with cooling chips efficiently.



Demonstration of how quadcore (top) speeds up a typical scientific calculation.

The four processors in a quadcore system are practically identical to the dual and single core processors on the market today, allowing them to run

several applications in parallel with big boosts in overall performance. This was illustrated at the launch event through a demonstration of how quadcore can speed up the rate of a typical scientific calculation involving a parallelized version of the program ROOT, widely used in the high-energy physics community, by nearly a factor of four over conventional single core processors (see image). Examples of similar gains from the fields of image rendering and computer games were also described at the event.

The multicore approach allows Intel to pursue the famous “Moore’s law”, named after Gordon Moore, one of Intel’s founders. This law – actually an observation of technological trends – states that the number of transistors on chips doubles roughly every two years. This results in a similar growth rate for processor power. Intel representatives are optimistic that the progress characterized by Moore’s law can continue for at least another decade.

François Grey, IT/DI

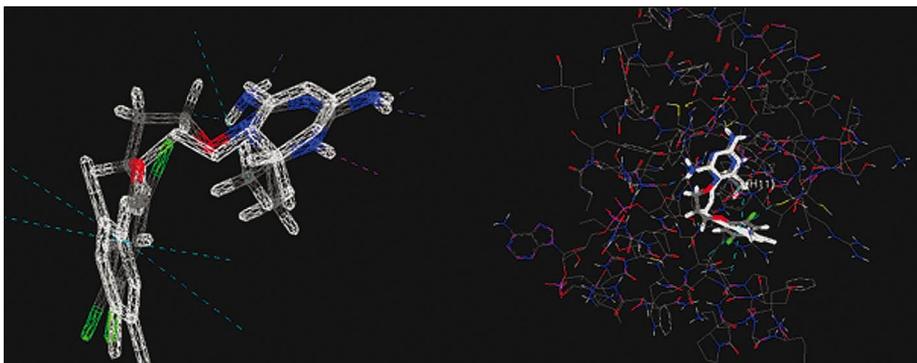
WISDOM continues the battle against malaria

In the often poorly attended fight against neglected diseases, the Worldwide in Silico Docking on Malaria initiative is proving to be a strong contender. WISDOM, which launched its second assault on malaria last month, harnesses the power of Grid computing to search for drug candidates to treat the disease.

“Drug development is a very expensive process,” explained Vincent Breton, WISDOM coordinator from LPC Clermont-Ferrand, “but by using Grid technology and with the support of our partners in academia and industry we are able to significantly cut costs for its initial step. By finding good drug candidates for neglected diseases in such a rapid and cost-effective way we can help to move away from the sad economic situation that made them neglected in the first place.”

WISDOM’s first bout against malaria in the summer of 2005 simulated the docking of 41 million potential drugs against target proteins on the malaria parasite, using the equivalent of 80 years of CPU time in six weeks. The parasite inhibits the breakdown of human haemoglobin, the protein that transports oxygen within the human body. The first WISDOM study identified three molecules that inhibit the haemoglobin breakdown, molecules which will now be tested further by laboratories.

This strategy demonstrated how Grid computing can be used to accelerate drug discovery research by speeding up the virtual screening process and reducing the cost of developing new drugs.



Left: structure of a potential anti-malarial drug. Right: a simulation of the drug binding to a protein from a malaria parasite. (Courtesy Vinod KASAM, CNRS.IN2P3.)

After searching for drugs to combat avian flu earlier this year the WISDOM initiative has returned to the search for potential malaria drugs. The most recent study, which ran until 15 December 2006, included participation from additional research groups around the world. Work with these new groups generated four target proteins for the study, two of which were previously unknown.

The search for potential drugs uses the FlexX commercial docking program, provided by free license from the German company BioSolveIT.

“BioSolveIT is happy to sponsor WISDOM,” said BioSolveIT CEO, Christian Lemmen. “The initiative takes full advantage of the speed and accuracy of FlexX – demonstrating the impact of virtual

screening in the search for drugs against neglected diseases.”

With the support of Enabling Grids for E-science and several related European Grid projects, WISDOM is currently testing up to 150 000 docked compounds per hour on 3000 computers around the world and will begin to report on the results of the work in early 2007.

For more information and a list of WISDOM partners visit the initiative’s website at <http://wisdom.healthgrid.org/>. More information about BioSolveIT and FlexX is available at the company’s website at www.biosolveit.de/.

● This article is extracted from the *iSGTW Newsletter* issue 16 November 2006 (www.isgtw.org/?pid=1000127).

Nathanael Verhaeghe, HealthGrid

IT launches Gridcast at Supercomputing 2006

CERN was present for the first time on a stand at Supercomputing 2006 (SC’06), the premiere US conference and trade show on high performance computing. SC’06 took place in Tampa, Florida in November with some 7000 participants. The California Institute of Technology (Caltech) graciously shared its stand with CERN, and four representatives of the IT department and the LCG project presented the latest results concerning Grid monitoring and Grid interoperability at the stand. The Enabling Grids for E-science project also had a booth at the conference.

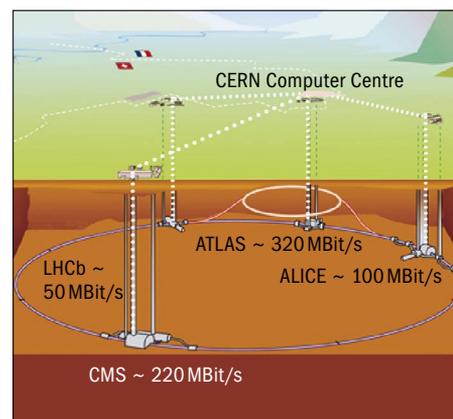
A novel feature was a “Gridcast” from the conference, consisting of a shared blog from CERN representatives on the stand and a series of podcast interviews with people at the conference (www.gridcast.org). This proved popular with people who were not able to attend the meeting as a way to get a better sense of what was going on, not just on the Caltech–CERN stand, but



CERN participants at the Caltech–CERN stand at SC’06 in Tampa, Florida.

at the event in general.

A Flash animation illustrating the life-cycle of data from the LHC proved popular with visitors to the stand and can be downloaded on the GridCafé website



A representation of the life cycle of LHC data.

at www.gridcafe.org. HP and Intel co-sponsored the Caltech–CERN stand, and Caltech achieved some breakthrough results in its data challenge activities.

IT Communications Team, DI, CERN

The Grid improves efficiency of Geant4 production cycles

Geant4 is a toolkit to simulate particle interactions with matter (Agostini *et al.* 2003). The simulated particles are propagated through magnetic and electrical fields, and through the materials of the detectors. It is used in high-energy physics experiments (BaBar, HARP, ATLAS, CMS, LHCb, etc), as well as in other fields such as space science, medical physics and radiation studies (Pia and Knobloch 2002, Allisson *et al.* 2006).

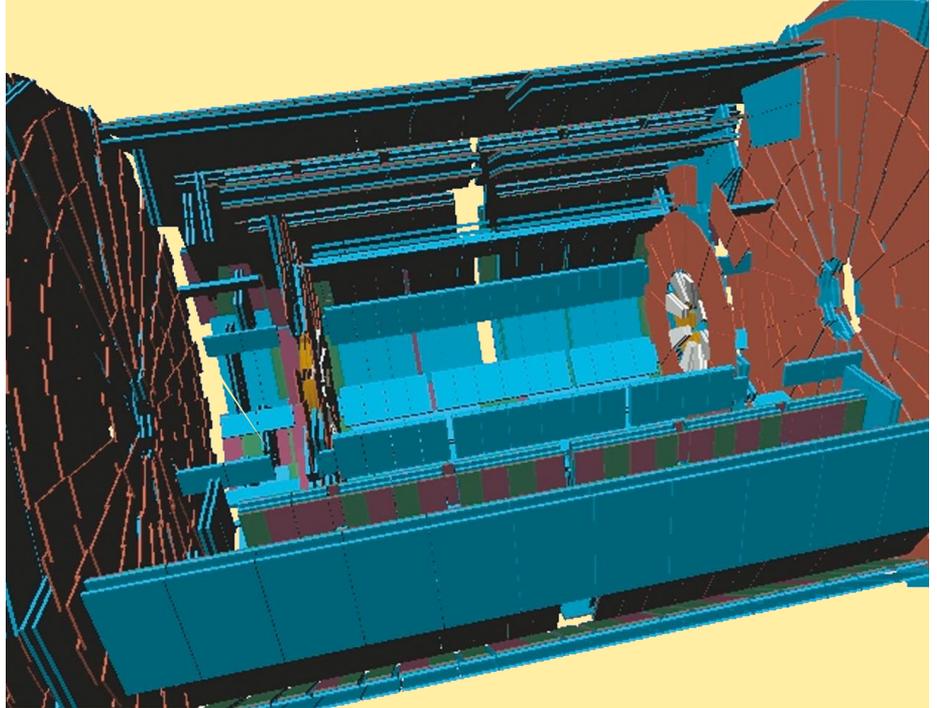
Geant4 consists of roughly half a million lines of C++ code containing components to model in detail the geometry and materials of complex particle detectors, in addition to a comprehensive set of electromagnetic and hadronic physics processes. Given its complexity and its wide range of applications it is mandatory to carefully and thoroughly test and validate each of its components, especially before major releases, generally twice a year.

During the testing phase we perform a set of regression tests, comparing two Geant4 versions to detect any significant difference between them by looking at calorimeter observables. The idea is to check automatically, using a variety of statistical tests, a large number of distributions. For those that are statistically different we produce a figure showing the two distributions in the same plot. Finally, this set of figures needs to be visually examined by a Geant4 expert in order to understand the cause of such difference, either a statistical fluctuation or a bug or a change in the physical models.

A test suite consisting of a simple, configurable, cylindrical sampling calorimeter aims to reproduce, in a simplified way (without the geometrical details or the instrumental effects of a real calorimeter), the major calorimeter types used in the LHC experiments: ion–scintillator; copper–scintillator; copper–liquid argon and tungsten–liquid argon (LHC hadronic calorimeters); lead–scintillator; lead–liquid argon; and PbW04 (LHC electromagnetic calorimeters).

As particles enter the calorimeter, we consider positive and negative pions; positive, negative and neutral kaons; and protons, neutrons and electrons.

Beam energies range from 1 GeV to 300 GeV – 23 energies in total. The motivation for having so many values is that showers are largely determined by the very first interaction. In addition there are several physics models available in Geant4



A view of the ATLAS detector muon system, simulated in Geant4. (Courtesy ATLAS collaboration. Thanks to J Apostolakis and A Ribon from PH/SFT.)

(during the regression tests we test five different physics lists, each of them being a consistent set of five physics models) that can be activated only within a certain kinetic energy interval of the incoming particle.

The observables considered are the visible energy deposited in each active layer, the total energy deposited in the whole calorimeter, the energy deposited in each active layer and the energy deposited in each radial ring.

Each parameter combination (energy–particle–material–physics list) defines a job and each job will produce 5000 events for the reference Geant4 version and 5000 events for the candidate one. Each set of events is stored in a separate n-tuple. Then a stand-alone program reads the two n-tuples and performs the statistical tests between all the pairs of observable distributions.

Geant4 in the Grid

Since December 2004, the regression tests of Geant4 have been executed in the Grid environment. To test the various hadronic models in a large number of configurations of beam particle type, beam energy, materials and physics configurations we need a few years of CPU time. All of this

CPU time is needed during a short period, about two weeks before a Geant4 release.

During the first production of the Worldwide LHC Computing Grid (WLCG; <http://cern.ch/lcg>) team created an initial framework that basically used the Grid in order to submit, keep track of and retrieve the outputs of large bunches of jobs. The main purpose of this tool was to hide much of the complexity of the Grid environment and to involve the new users in an easy, fast and Grid-unaware way.

During the first week of production we distribute the Geant4 software (also containing the reference version) to all participating sites. The software is provided in a tar file format. All sites require locally 5 Gbyte of mirrored disk space in a file shared system, accessible by all worker nodes. During the software installation phase we distribute this tar file into the shared file system at all sites and perform a small production to validate this installation. Thanks to this small production we can select the sites that will run the production. In addition we can aid those sites having any configuration problems by debugging their systems.

During the second week we perform the

whole production, sending bunches of jobs – one for each physics list. In this phase the Geant4 developers continuously check the results and fix any eventual problems in the new candidate.

After the first production with Geant4 on the Grid the production tool was extended and generalized for any community or experiment. Since December 2005 the Geant4 collaboration has been a fully recognized virtual organization in the WLCG/Enabling Grids for E-science environment (EGEE; <http://public.eu-egee.org>). Several sites have provided the required resources, guaranteeing the success of the production and achieving an efficiency of 100% during the last two productions. Within the Geant4 production we define efficiency by the number of submitted jobs compared with the number of successful jobs returning the desired output.

In terms of resources and services, Geant4 asks for a minimum of 120 CPUs for the whole production. Normally CERN provides 50% of the required CPUs. The rest is spread over 30 sites currently supporting the Geant4 collaboration.

Most of the Grid services are also provided by and centralized at CERN. The rest of the sites ensure access to their batch system with dedicated queues for Geant4 and access to the shared file system for software installation purposes.

The total output of the production, about

20 Gbyte, is fully retrieved in an AFS area of Geant4 created for this purpose.

Optimization of the Geant4 production

In June 2006 we included the Ganga/DIANE framework developed by the ARDA (A Realisation of Distributed Analysis) team at CERN, to increase the performance of the Geant4 production and to optimize the use of the Grid resources (Ganga: A Grid User Interface, <http://cern.ch/ganga>; DIANE: a Distributed Analysis Environment, <http://cern.ch/DIANE>; ARDA; <http://cern.ch/arda>).

Ganga is a tool jointly developed for the ATLAS and LHCb experiments at CERN. Ganga simplifies access to the Grid for the end users and is a convenient job management tool. The history of the user jobs is stored in a personal job repository; the status of submitted and running jobs is monitored and the output automatically retrieved. Ganga provides an extensible framework that supports various applications and helps to configure them, effectively hiding any differences between running the application locally or on the Grid. A user may interact with Ganga via a command line or a graphical user interface. Python is used as the interface language. Ganga is a tool suitable for running analysis jobs in high-energy physics but may also be used for any user Grid activities.

DIANE is a Grid optimization layer. When the user jobs are submitted to the

Grid, DIANE creates an overlay “virtual cluster” using available worker nodes. The cluster consists of a “master” agent and a number of “worker” agents. User input is split into a large number of fine-grained tasks. The execution of the tasks is coordinated by the master agent providing automatic load balancing and instant error recovery. The task execution flow and error recovery policies may be easily customized according to the specific application needs.

DIANE improves the job turnaround time and reliability of the application execution. It also enables estimates of the job completion time based on the delivered partial output.

The Ganga/DIANE framework greatly improved the performance of the Geant4 production. In addition, the run completion time is more predictable and allows a better planning of the production and less intervention from the production operator.

Further reading

S Agostini *et al.* 2003 Geant4: a simulation toolkit, *NIMA* **506** 250-303

J Allison *et al.* 2006 Geant4 developments and applications, *IEEE Transactions on Nuclear Science* **53** 270–278

M G Pia and J Knobloch 2002 Particle physics software aids space and medicine *CERN Courier* **42 (5)** 33; www.cerncourier.com/main/article/42/5/17/1

Patricia Mendez Lorenzo, CERN, IT/PSS

CERN gLite Grid community uses virtualization

These days we hear a lot about virtualization and its various benefits, such as improved security and the efficient utilization of processor cycles. In the IT/GD group we have been progressively using virtualization to manage our certification and development test bed for gLite middleware.

Virtualization is not a new technology and was used in various forms in mainframe computers. Today’s virtualization technology, however, is a layer of software that sits between the hardware and the host operating system and allows multiple operating systems to be run in parallel on a single machine. A number of virtual machine monitors are available, of which Xen and VMWare are two well-known examples. Virtualization provides a secure and isolated environment for all sorts of legacy software that can be run unmodified as virtual machines.

Our experience

In our test bed we have been running Xen with SLC4 machines. Our guest virtual machines have been both SLC3 and SLC4. Unfortunately, currently no open source tools exist that completely manage the deployment and management life-cycle

of virtual machines on a cluster. We have developed some proof-of-concept tools based on an HP product called Smart Framework for Object Groups (SmartFrog), which is a Java-based deployment framework for distributed components. Our tool stack, which we call VirtualGrid (vGrid), also includes higher level web interfaces written in Python and PHP, which wrap SmartFrog components to manage, monitor and deploy virtual machines on our test bed. This work was done in collaboration with the CERN openlab. We are now in the process of releasing these tools to the wider gLite community.

Under virtualization the atomic software entity is an operating system image, which is pre-configured with Xen and other user-defined software. In our case it is the gLite middleware. Once the virtual machine is deployed, from a user’s perspective it appears as a normal physical machine. The virtual machine image can be saved in a number of formats and can easily be redeployed elsewhere.

From this experience virtualization has greatly improved the efficiency of our software development life-cycle. During the middleware testing process, the entire test bed may have to be reinstalled from

scratch many times, which previously took a great deal of time, with each reinstallation measured in days. Now, using pre-configured virtual machines, the cycle of scrapping the existing virtual machines and redeploying the entire virtual certification test bed takes only 20–30 minutes.

Conclusion

Virtualization technology can not only play a vital role in improving productivity in the software life-cycle, as in this example, but it also holds an important place in the LHC Grid and the wider grid community by offering the possibility of providing secure, isolated and pre-configured execution environments (operating system with required compilers, libraries and applications) for Grid jobs.

Meanwhile the new generations of processor families from Intel and AMD, respectively the VT and Pacifica brands, have built-in hardware virtualization. These are expected to improve performance significantly by removing some software overhead. Such developments are very encouraging and give us a glimpse of the future, where virtualization technology will be widely used in all computer platforms.

Omer Khalid, IT/GD

CERN CA provides security

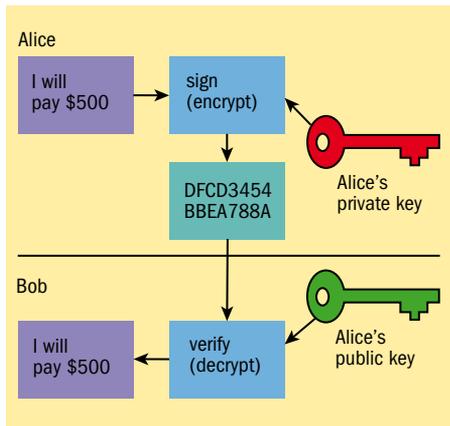
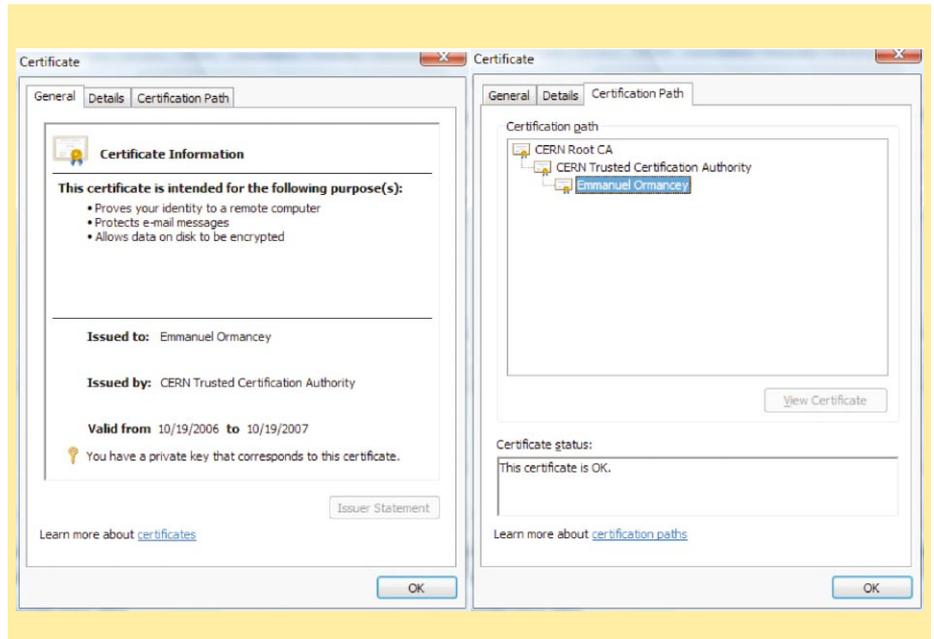


Fig. 1 (top). Only the person who possesses the private key can sign a message, while anyone can check the signature using the public key.

Fig. 2 (right). Authentication windows showing certificate information and the certification path. Users can access this information to verify authenticity.



Digital certificates are electronic credentials that are used to certify the identities of individuals, computers and other entities on a network.

Digital certificates function similarly to personal identification such as passports and ID cards. For example, passports and ID cards are issued by recognized government authorities and digital certificates are issued by recognized certification authorities (CAs).

When someone requests a passport or ID card the government authority verifies the identity of the requester, certifies that the requester meets all requirements to receive the card and then issues it. Before a certificate can be issued a CA or CA administrator must verify the requester's identity, determine that they meet all requirements to receive the certificate and then issue it.

Like personal identification such as a passport or ID card, a digital certificate can be used to verify the identity of its owner. When the digital certificate is presented to others it verifies the identity of its owner based on the quality of the contents of the certificate, which includes the following:

- Personal information that helps to identify the owner.
- The signature of the issuing authority. For digital certificates the issuing authority is the CA.
- Information needed to identify and contact the issuing authority.

In addition, the quality of a certificate is enhanced if it:

- Is designed to be tamper-resistant and difficult to counterfeit.
- Is issued by an authority that can revoke the certificate at any time (for example, if the employee to whom the certificate was issued is no longer employed by the firm).
- Can be checked for revocation by contacting the issuing authority.

Public key cryptography

A certification authority is based on a public key infrastructure. Public key cryptography is a form of cryptography that generally allows users to communicate securely without having prior access to a shared secret key. This is done by using a pair of cryptographic keys designated as a public key and a private key, which are mathematically related.

The private key is kept secret, while the public key may be widely distributed. In a sense, one key “locks” a lock while the other is required to “unlock” it. It should not be feasible to deduce the private key of a pair given the public key, and in high quality algorithms no such technique is known.

One analogy involves a locked store front door with a mail slot. It is exposed and accessible to the public; its location (the street address) is in essence the public key. Anyone knowing the address can drop a written message through the slot. However, only the person who owns the matching private key, the store owner in this case, can open the door and read the message.

Using a private key to encrypt (thus signing) a message, anyone can check the

signature using the public key. Validity depends on private key security (figure 1).

Use of certificates

CERN certificate services allow strong security based on public key encryption that can enhance a variety of internal and external applications, including authentication, encryption and data integrity.

For authentication, digital certificates help verify identity because the data in a certificate includes the public cryptographic key from the certificate subject's public and private key pair. A message signed with its sender's private key can be verified by the recipient of the message as authentic by using the sender's public key, which can be found on a copy of the sender's certificate (figure 2). Verifying a signature by using a public key from a certificate proves that the signature was produced using the private key of the certificate subject.

Encryption certificates enable privacy for data that is transmitted using a number of different methods. Some of the more common privacy-enabling protocols that use certificates are secure multipurpose Internet mail extensions (S/MIME), transport layer security (TLS) and encrypting file system (EFS).

For data integrity, digital signatures are used when data is distributed in plain text or unencrypted form. In these cases, although the sensitivity of the message itself might not warrant encryption, there could be a compelling reason to ensure that the data is in its original form and has not



Fig. 3. Steps in the process of requesting a user certificate. Anti-clockwise from top left: Browser or manual operation selection; date of birth validation; request using Internet Explorer; and request using OpenSSL textbox. See also <http://cern.ch/ca>.

been sent by an impostor.

CERN Certification Authority

CERN is now running a certification authority (CERN CA) as an official service, replacing the legacy certification authority ran by LCG in an offline mode, dedicated for Grid users. The new CERN CA can provide certificates to all CERN users as well as CERN hosts. This new service delivers certificates in real time: when a user requests a certificate, identity is verified within the CERN human resources (HR) database and the certificate can be downloaded instantly.

The certificates issued by CERN CA are automatically trusted inside CERN and within the Grid community. CERN CA certificates are approved by EuGridPMA (www.eugridpma.org) and can be used to authenticate on Grid applications. Outside CERN, trusting those certificates can be achieved in two clicks using a web browser (figure 2).

Who can request a certificate?

To request a user certificate a user must:

- Be registered in CERN's HR database under one of the following status categories (for which physical presence at the appropriate registration service is required): STAFF, FELL, PDAS, PJAS, USAS, CASS, UPAS, USER, DOCT, TECH, ADMI, SUMM, CHIL, APPR; or
- Be an employee of a CERN contractor (status: ENTC).
- Have a CERN computer account and register an e-mail address. The user's group manager creates these manually.

Only users registered in CERN's central network database (LANDB) as the computer administrator can request a host certificate

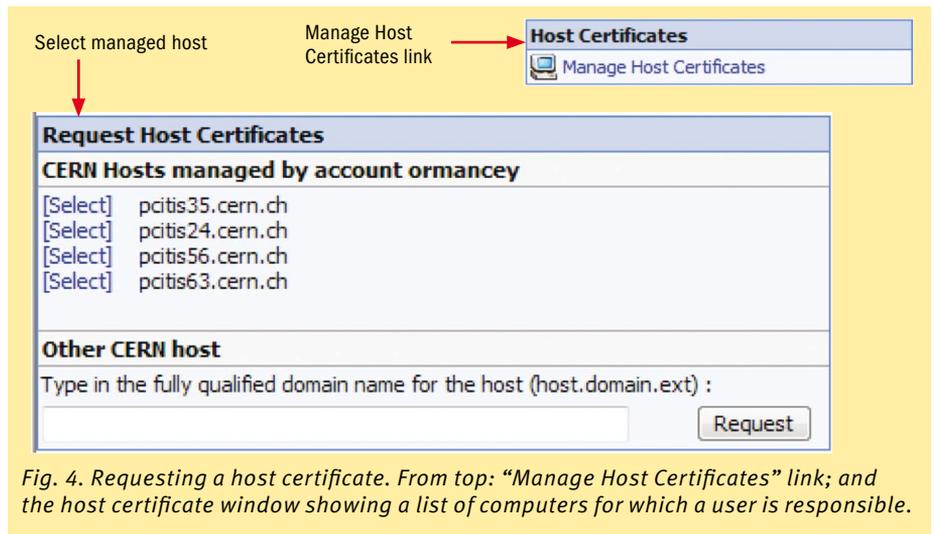


Fig. 4. Requesting a host certificate. From top: "Manage Host Certificates" link; and the host certificate window showing a list of computers for which a user is responsible.

for the computers they manage, and they must already have a user certificate. If a user does not meet these requirements, he/she will have to request a certificate from his/her home institute if necessary.

How to request a user certificate

Requesting a user certificate is an easy operation that can be achieved using the CERN CA web interface (<http://cern.ch/ca>). The interface supports automatic requests using Internet Explorer and Mozilla/Firefox, as well as a manual request form for OpenSSL-generated requests. Please note that in addition to your CERN credentials, your date of birth will be used for identity validation (figure 3).

How to request a host certificate

Requesting a host certificate is also a straightforward procedure, but a couple of

requirements must be met: the user must have a valid user certificate and this must be authenticated via <https://cern.ch/ca>.

If this is in place the user can simply click on the "Manage Host Certificates" link on the CERN CA website. A list will then be displayed detailing the computers for which the user is declared either "responsible user" or "main user" in LANDB on <https://network.cern.ch>. Simply click on the desired host to start the "Request Certificate" procedure (figure 4).

You cannot request a certificate for a host for which you are not responsible or for which you are not the main user.

What to do with a certificate

If you are a Grid user you need a certificate to authenticate and access LCG applications. Classic CERN users can take advantage of certificates to simplify the

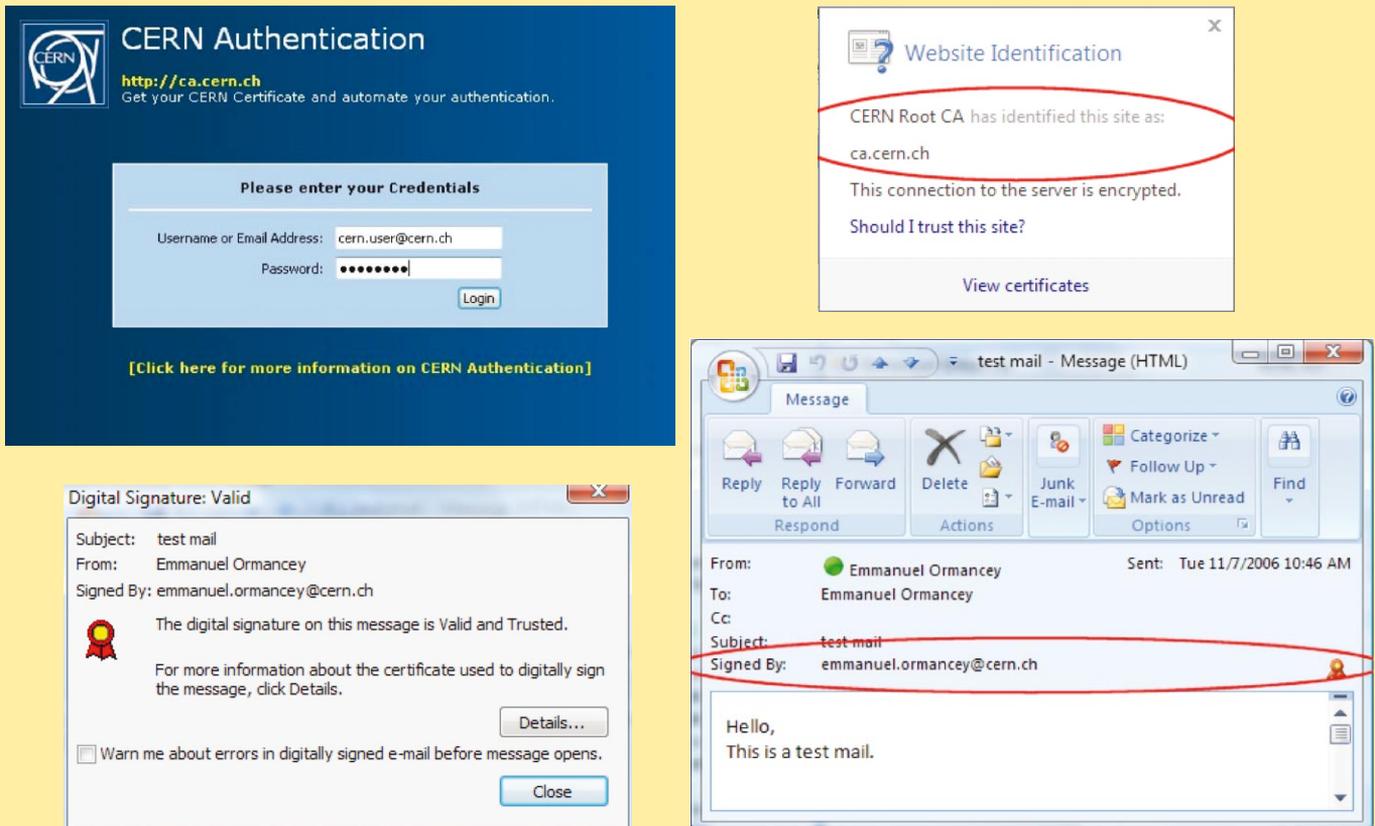


Fig. 5. Certificate uses, clockwise from top left: CERN authentication form, which allows simplification of the login process; certificate check on SSL website; signed mail enables authentication of the sender's identity; and signature verification on an e-mail.



Fig. 6. The future of CERN certificates is SmartCards. From left to right: WindowsXP logon window with a card reader detected; when a SmartCard is inserted Windows requests the PIN; and CERN access card with SmartCard chip.

login process: with a certificate it is not necessary to type in credentials.

The goal is to simplify the CERN authentication on various applications, such as EDH, AIS applications, Remedy, EDMS and various central services. A common authentication interface will soon be activated and will support user certificates as well as classic login and password authentication.

User certificates can also be used to sign mails as verification for recipients that the message originated from the sender named and was not modified during transport (figure 5).

Host certificates can be used to secure communications between hosts or between

a host and a client. A typical use would be enabling SSL on a Web server to allow https mode. The certificate proves that you are on the correct site (figure 5).

The future of CERN certificates

CERN has already taken a major step forward in authentication and identity validation in computing resources by establishing its own certification authority. However, by integrating SmartCards into the process CERN could further develop its use of certificates.

A SmartCard is an electronic chip that stores and secures the user certificate. A certificate stored on a SmartCard can be used for classic certificate authentication

schemes, as well as for desktop authentication. Instead of typing in your credentials you would simply insert your SmartCard into a card reader to login or unlock your desktop computer. For example in Windows, the login process allows SmartCards as soon as a card reader is detected (figure 6).

The idea is to integrate the SmartCard with the CERN access card to have all access controls on one device. Currently a SmartCard pilot is running and has generated interest in many areas. For example the ALICE experiment plans to request SmartCards to access the control stations and increase security.

Emmanuel Ormancey (IT/IS)

HENP community shares technology at HEPiX meeting

HEPiX is a global organization of service managers and support staff that provides computing facilities for the high-energy and nuclear physics (HENP) community. All operating systems that HENP uses are covered including Unix, Linux and Windows. HEPiX holds regular meetings for participants to present recent work and future plans, and to share experiences.

The autumn 2006 HEPiX meeting was held in October at Jefferson Lab, a nuclear-physics laboratory in Newport News, Virginia, some 250 km south-east of Washington DC. It was the third time the meeting has been held there. Attendance at the meeting was rather small (compared with the last US meeting in autumn 2005) with fewer than 50 participants – half the number of participants of the past two European meetings.

Roy Whitney, the chief information officer of the newly reorganized IT department at Jefferson Lab, opened the meeting by describing the mission and activities of the laboratory, which centre on nuclear physics. He described some of the interesting physics from the main experimental halls and added that Jefferson Lab plans to upgrade the accelerator to 12 GeV.

Meeting Highlights

Collaboration was a theme of the week with:

- A suggestion to create a HEPiX e-mail whitelist (Jim Fromm from Fermilab spoke about SpamCop). Kevin Hill at FNAL has taken the first steps to set up this new list.
- A suggestion that HEPiX assist LCG to improve and standardize Tier-1 and Tier-2 site monitoring (Ian Bird from CERN presented Site manageability issues for LCG). LCG proposes a site-availability measurement based on the availability of a defined set of core services. These data should be further used to trigger alarms when the metric drops below a threshold.
- Randy Sobie, chair of the International High Energy Physics Computing Coordination Committee (IHEPCCC) spoke via an audio link from Victoria, British Columbia, Canada, to propose creating joint study groups on benchmarking and on global file systems. IHEPCCC proposes the Requirements Technical Assessment Group (RTAG) model, which the LCG used successfully.

As usual security was another hot topic with the introduction of two factor authorization issues mentioned by at least two major laboratories, Brookhaven National Laboratory (BNL) and Jefferson Lab; two spamming war stories; and a



The audience in the meeting room at HEPiX 2006, held at Jefferson Lab in October.

very detailed tale of a particularly serious attack lasting around 12 months: the Stakkato Intrusion, a security incident at the National Supercomputer Centre (NSC) in Sweden, which started in 2004 and culminated in the summer of 2005 with the arrest of a suspect. Leif Nixon from NSC estimated that over a thousand sites all over the world were affected at one time or another during the attack.

Fred Dylla, the Jefferson Lab chief technology officer, gave an interesting and entertaining talk at the Conference Banquet on the importance of vacuum tubes over the years (“From a spark in a vacuum to sparking the vacuum”). The speaker covered 200 years of the history of vacuum tubes in 20 minutes. This was probably the first such after-dinner talk for HEPiX and was definitely appreciated by the audience.

A number of sites (e.g. INFN, DAPNIA, SLAC) reported on tests conducted and/or planned on global file systems such as Lustre (Cluster File Systems, Inc.) and General Parallel File System (GPFS, IBM). GPFS has shown enough robustness and reliability that it is being adopted by more and more sites for shared file systems. Also, Fermilab is investigating a new device called BlueArc. They use BlueArc network attached storage (NAS) devices for online storage systems and are very pleased with its performance (although it is still quite expensive).

Yet again Michel Jouvin of Laboratoire de l'Accélérateur Linéaire (LAL) has won the (virtual) award for the HEPiX member adopting most CERN software, discovering the CERN-dependencies where they remain, and spreading the word that sharing is

a “good thing”. Next on his list, after he has completed the switch to Indico, is the CERN IT/FIO's SLS service-level monitoring (Sebastian Lopienski presented the SLS project at CERN).

Trip report and online information

As usual those readers who would like to get more details can read the trip report written by Alan Silverman, CERN IT, which is available at www.hepix.org/mtg/fall_06_jlab/HEPiX%20_Lab_Trip_Report_silverman.pdf. This long report (17 pages) is mostly a personal set of notes but includes some input kindly supplied by Helge Meinhard and Sebastian Lopienski.

Readers interested in particular points are referred to the overheads also available at <http://conferences.jlab.org/HEPiX/> and at <https://indico.fnal.gov/conferenceDisplay.py?confid=384> under the link “Program”.

Future HEPiX meetings

The next meetings will take place:

- At DESY in Hamburg for the spring 2007 meeting (23–27 April)
- At the University of Washington Genome Sequencing Center, St Louis, for the autumn 2007 meeting, 5–9 November in the week preceding Supercomputing'07 in Reno, Nevada.
- At CERN for the spring 2008 meeting.

For the DESY meeting, topics suggested were benchmarking, cluster file systems and voice-over-Internet protocol (VoIP). In general, discussion topics (as opposed to LCG workshops) are likely to attract LCG Tier-2 sites.

Alan Silverman IT/DI and Nicole Crémel IT/UDS

CSC2006 sets three new records



Clockwise from top: CSC2006 participants; a morning lecture in the auditorium at the Kumpula Campus; and the Helsinki harbour.

Though it is not a goal in itself, the CERN School of Computing (CSC) seems destined to break records every year. Last year it broke the record for the number of nationalities (27) of the participants (79 in total).

CSC2006, which took place at the Kumpula Campus in Helsinki from 21 August to 1 September 2006, broke three records: the number of nationalities (now up to 29), the proportion of females (24%) and the number of students who registered for the final optional exam (totalling 75 out of 79 participants).

The three new records, which may be hard to beat in the future, signify the evolution of the school over the past few years. First, the involvement of participants from a growing number of countries shows the increasing universality of the CSCs. This is largely due to the support of the European Commission, which provides grants to students through a four-year programme under the Marie Curie Actions. These grants cover all costs for some 20 young engineers and scientists coming from the world's less favoured regions. This year, in addition to participants from CERN member states, students came from Algeria, Brazil, China, Columbia, India, Iran, Lithuania, Mexico, Romania, Russia,

Serbia and Montenegro, Turkey and the US.

The second record reflects the importance that the CSC places on equal opportunities. Female engineers and scientists are strongly encouraged to apply, and those selected to attend are in turn invited to encourage their female colleagues to apply to future schools. The director of the CSC has appointed Wisla Carena to act as equal opportunity officer for the school, to analyse the participation, understand difficulties or obstacles that female engineers may face, and take any action to keep increasing their participation.

The third record is a sign of the increase in student participation, an effort the school initiated three years ago to encourage students to have a more active attitude. The final examination is one component of a series of actions undertaken in that direction, and the ever-increasing interest by participants shows their willingness to be active. Other actions include a contest for students to present topics of their choice, hands-on mini-projects leading to oral reporting, as well as preparatory discussions for the inverted CSC.

In this spirit, one novelty of the year was the organization of a full sports programme for those interested. Benefiting from the

proximity of a well-equipped sports centre, activities ranged every day from basketball, volleyball, floorball, badminton, climbing lessons or swimming. The now traditional mixed-doubles badminton tournament and the end-of-school football match, played this year next to the famous Helsinki Olympic Stadium, complemented the new activities.

CSCs are also about partnership and collaboration: every school is co-organized with one or more national scientific institutions in the host country. This year the school was organized with the Helsinki Institute of Physics (HIP). With the full support of the HIP director, Prof. Dan-Olof Riska, a dedicated team, coordinated by Veikko Karimaki and Jorma Tuominiemi, seconded the CERN organizers, taking the lead in several domains, such as the sport and after-dinner programmes. Special thanks go to Tapio Lampen and Aatos Heikkinen, who became the stars of the school for leading these two special programmes.

CSC2006 is now over but will remain a memorable experience for those who attended. Croatia will host CSC2007.

For more information on CSC see <http://cern.ch/csc>.

François Fluckiger, IT/DI

TWiki facilitates collaborative information handling at CERN

TWiki is a software collaboration tool that was introduced at CERN in 2003 and has been growing in popularity ever since. It now has more than 1000 active editors and 15000 documents. Pete Jones of IT/DES and a member of the TWiki support team talks to CNL about the tool. He explains why it has become so popular and describes how users can create and edit documents with the tool.

What is TWiki?

TWiki is a variant of a Wiki. A Wiki is a piece of software that enables users to freely create and edit Web pages using any Web browser. A Wiki website looks like most others, but it has an edit button on every page. The best known example is *wikipedia.org*, an online encyclopedia-building collaboration.

There are many different types of Wiki. Examples are JaWiki (a small Java servlet-based Wiki), phpwiki (a php-based Wiki) and Confluence (a commercial Wiki written in Java). CERN IT's chosen solution is called TWiki. This is a Perl-based Wiki that was developed by Take Five and is GNU general public licensed software.

Who uses TWiki?

TWiki is installed on many corporate websites, including those of British Telecom, Michelin, Motorola, Disney and Yahoo. At CERN it is widely used by the LHC experiments, the Grid and by the IT department.

Why use TWiki?

TWiki enables a group of users to work together to create and edit documents for many different purposes including:

- designing and documenting software projects;
- developing a knowledge base and FAQ system;
- scheduling events by using calendar features;
- operating an internal message board;
- tracking issues (i.e. bugs) and features;
- managing documents;

- archiving software; and
- writing and storing minutes of meetings.

How does one use TWiki at CERN?

Go to <https://cern.ch/twiki> using a Web browser. The left-hand menu lists each of the existing Web collaborations, and clicking on one of these will take you to the documents of the collaboration (some pages might not be available to you because access control has been set).

Users need authentication to edit a document – at the moment this is based on the user's AFS password. Two editors are available for each page: a simple text-based one that understands HTML syntax and TWiki's own simple formatting tags; and a new WYSIWYG editor that enables documents to be edited without the need to type or remember formatting tags.

If anyone can edit a document won't this cause problems?

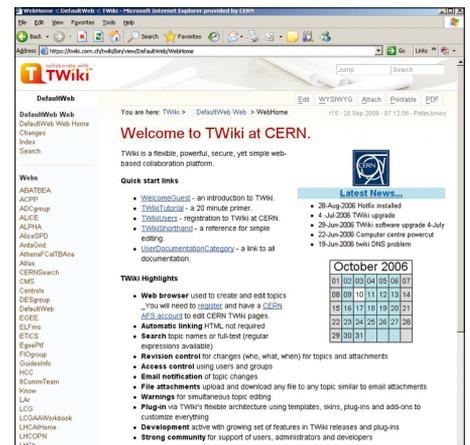
Open freeform editing is the essence of TWiki culture. However, for the CERN TWiki, the editor must be a registered CERN user. Peer influence also ensures that topic content is relevant and focused. TWiki uses a revision control system that stores and logs all changes made to a topic so that one can revert to a previous version of a document. If necessary, a topic or even a whole Web can be read-write protected by using access control settings.

Where does TWiki run and where are the documents stored?

TWiki runs on three dedicated dual-processor Linux PCs. These machines are load balanced to provide the fastest service and are located in the Computer Centre in B513. Each document created by TWiki is stored in an AFS directory. This directory is backed up every night and the data is also copied at regular intervals during the day. An offsite back-up is envisaged for the near future.

Is there support for this software?

TWiki was first released in 1998, it has had



The CERN TWiki website offers a tutorial and guides to document editing.

nine major version releases, and there are now more than 100 developers contributing to the TWiki code. Users at CERN can request support from twiki.support@cern.ch, and the *TWiki.org* website offers extensive documentation and online forums.

How can I start to use TWiki?

TWiki has a sandbox Web (available from <https://twiki.cern.ch/twiki>) where users can test out the software. Once you have a CERN AFS account, click on the "Sandbox" link found on the left menu and then create a new document by selecting either the "Create TestTopic" button or the "Create by Name" button. Simply type and save.

What are plug-ins?

These are enhancements to TWiki to create calendars, database access, drawings, spreadsheets, charts, graphs and more. In your test document type in `%CALENDAR%` to see how easy they are to use.

Where can I get more information?

A complete list of documentation can be found on <https://twiki.cern.ch/twiki/bin/view/TWiki/UserDocumentationCategory>. For further enquiries contact twiki.support@cern.ch.

If you want to be informed by e-mail when a new CNL is available subscribe to the mailing list cern-cn1-info. You can do this from the CERN CNL website at <http://cern.ch/cnl/>.

How to help keep your PC secure

Did you know that even if your PC has up-to-date patches, the latest anti-virus software and runs a local firewall, it can still be infected? When computers are used for personal rather than professional use the chance of infections and other security incidents rises – movies, games, music and other personal applications all have risks.

Here's some advice to help keep your PC secure. Although useful for all platforms, this advice is particularly targeted at Windows users.

- **Configure to run without administrator privileges:** There are a growing number of “zero-day exploits” – security weaknesses that are discovered before patches become available. With these exploits simply clicking a Web link while you have administrator privileges could automatically install malicious software on your machine. You are recommended to run without administrator privileges as this restricts the damage that malicious software can do. For information on running without administrator rights see: <http://cern.ch/WinServices/Help/?kbid=010121>.
- **Be cautious of pop-ups:** Some pop-ups may be configured so that even if you click “Cancel” or “No” or try to close the window a program could still be executed. On a Windows PC close the active window by pressing [Alt][F4].

- **Do not install plug-ins:** Installing a plug-in could also download any malicious software that the plug-in might contain. If a website requires a plug-in to view it, it is best to avoid using it.

- **Ensure software downloads respect copyright and licensing:** Not only for legal reasons but also because “free” versions of copyrighted software often contain Trojan horses, spyware or other malicious software that could infect a PC. Spyware is often included in “free” software and is used to trace your activity and possibly even the data you type, including passwords. For more information about spyware and how to avoid it, see <http://cern.ch/WinServices/Help/?fdid=16>.

- **Do not expose your password:** Never use a CERN password for private use and be cautious of attempts to “steal” your password. CERN’s computing staff, including the computer security team, will never ask for your password, so be wary of e-mails, instant messages and chat that request your password, including via Web links. This trick is known as “phishing” (password fishing). If you think your password may have been exposed, follow this link to change it: <http://cern.ch/security/passwords/>.

- **Trust CERN’s security solutions:** Some of the “free” security software advertised on the Internet can contain malicious

software, particularly “free” software for Windows PCs. This form of “social engineering” hides malicious software inside a security package to make you think that you can trust it. Don’t be tricked into installing it.

The above advice may help you avoid some of the more recent tricks used by attackers. In addition, continue to follow these standard recommendations to keep your computer secure:

- Do not click on Web links in unexpected e-mails, spam, instant messages and chat.
- Do not open unexpected attachments.
- Do not install additional software and plug-ins, even when the content seems to be an extension of legitimate software.
- Lock your screen each time you leave your office. From a Windows PC use [Control][Alt][Delete] and select “Lock Computer”, or if you have a Windows keyboard, simply press [Windows][L].
- Keep yourself informed of CERN’s computing rules (<http://cern.ch/ComputingRules/>) and security advice (<http://cern.ch/security>).

CERN has put in place a number of mechanisms to secure the site, but technical solutions cannot prevent all problems. When using a computer there are always risks, but the advice in this article can help to minimize them.

Kate Bradshaw and Denise Heagerty, IT/DI

CERN Computing Colloquia Spring/Summer series kicks off with open source software talk

The CERN Computing Colloquia Spring/Summer series in 2007 will focus on the topic open source software and will kick off early this year on 28 February with a talk about open source from an African perspective by Mark Shuttleworth, founder of the Ubuntu project, an enterprise Linux distribution recognized for its ease of use.

Mark Shuttleworth is an entrepreneur renowned in South Africa as the first “Afronaut” – he was the first African in space and the world’s second tourist cosmonaut. He has launched a range of projects, such as the Freedom Toaster, to help ease access to digital information

and open source software in Africa. His Shuttleworth Foundation supports initiatives in education, technology and open content in Africa.

Meanwhile, the 2006 Autumn/Winter Series on mobile communications will wrap up, somewhat belatedly, on Friday March 9th, with a talk by Sir Robin Saxby, chairman of the UK company ARM. If you have a mobile phone or other mobile device, you probably have a processor designed by ARM. Today, ARM processors account for more than 75% of all 32-bit embedded CPUs, a characteristic feature being their low electric power consumption.

Sir Robin will describe his vision of the future of mobile communications, and explain how a small British design company became a world leader in IT.

Sir Robin is currently president of the UK’s Institution of Engineering and Technology, a founder of the the Advanced Internet Methods and Emergent Systems Centre at the University of Liverpool, and a pioneer in the transfer of e-science and Grid activities to e-business.

For details of time, location and abstracts of the colloquia, visit <http://computing-colloquia.web.cern.ch/computing-colloquia/>.

The deadline for submissions to the next issue of the CNL is

2 March 2007

Please e-mail your contributions to cnl.editor@cern.ch.

openlab invites applications for its latest student programme

The CERN openlab student programme, which will celebrate its fifth anniversary in 2007, is open for applications from bachelor, Masters and PhD students in computing and physics.

Successful applicants will spend two months at CERN this summer working on cutting-edge Grid-related technology projects. They will receive training in all aspects of Grids from the



LCG, Enabling Grids for E-science (EGEE) and CERN openlab staff. Last year the programme, which was co-sponsored by HP and Intel, accepted some 20 students. See www.cern.ch/openlab for more information.

Candidates should send a CV and letter of support from a supervisor to Francois.Grey@cern.ch. The closing date for applications is 31 March.

François Grey, IT/DI

NICE streamlines printing infrastructure

NICE have removed the Windows CERN Printer Wizard and Unix "xprint" command. The new printing system is in full production and hosts all centrally managed printers and plotters at CERN. Since 11 December 2006:

- the Windows CERN Printer Wizard has been uninstalled on all centrally managed NICE computers; and
- the "xprint" command is no longer on SLC3 and SLC4 Linux computers at CERN.

We remind users that full documentation on printing at CERN is available at <http://cern.ch/WinServices/Help/?kbid=070001>.

The NICE team

Solaris service stops

As announced in April 2005 to the desktop forum and in a *CNL* article (see <http://cerncourier.com/articles/cnl/2/6/6/1>), Solaris is no longer an LHC Physics platform. The SUNDEV service at CERN ceased on 1 July 2005.

In line with the IT resource constraints and security policy, Solaris User and OS Support has been stopped from 2007. Service managers have been informed and advice was given on a possible migration to another IT-supported platform.

This means that, as from 1 January 2007, support will no longer be provided by the helpdesk or IT staff. E-mails to Solaris.Support@cern.ch will receive an automatic reply containing this statement.

However, CERN users can visit SUN Microsystems Support at http://ch.sun.com/ch.sun.com/support/index_en.jsp to check if SUN provides the kind of Solaris support they need.

The User Support team

Recent changes to IT services

Changes to services in the IT department are published on the service status board (SSB), which is located at <http://cern.ch/it-servicestatus>. The SSB also

includes services incidents, scheduled interventions, power-cuts and the status of most services. The most recent changes and their dates of posting are shown below.

As of:	Service change:
29 January 2007	SLC3 to SLC4 migration for LXPLUS and LXBATCH alias switch delay
16 January 2007	Firewall BLOCK affecting Symantec, UDP/TCP port 2967
09 January 2007	MAC-address authentication on the Technical Network
07 January 2007	DEVDB10 upgraded to 10.2.0.3
06 December 2006	Removal of Windows CERN Printer Wizard and xprint
27 November 2006	There will be some restrictions on updates in the Network Register Web form
06 November 2006	Access restrictions enforced for Technical Network (TN)
30 October 2006	Deployment of Castor client software version 2.1.1-4
02 October 2006	Closure of insecure pop, imap and smtp access to central e-mail servers
01 October 2006	castorgrid.cern.ch SRM endpoint to be removed
25 September 2006	Newsgroups service relocated to new server

Calendar

January

29–2 Feb **OGF19: Open Grid Forum**
Chapel Hill, NC, US
www.ogf.org

February

15–19 **AAAS: American Academy of Arts and Sciences annual meeting**
San Francisco, CA, US
www.aaas.org/meetings/Annual_Meeting/

March

1–2 **Spanish Conference on e-Science Grid Computing 2007**
Madrid, Spain
<http://webrrt.ciemat.es:8000/e-science/index.html>

3–6 **WEBIST 2007: 3rd International Conference on Web Information Systems and Technologies**

Barcelona, Spain
www.webist.org

26–30 **IPDPS 2007: 21st IEEE International Parallel & Distributed Processing Symposium**

CALL FOR PAPERS
Long Beach, CA, US
www.ipdps.org

April **DEISA Symposium 2007**

Munich, Germany
www.deisa.org/news_events/future_events.php

23–27 **HEP*X* 2007 spring meeting**
Hamburg, Germany

www.hepix.org

23–27 **ACAT 2007: Advanced Computing and Analysis Techniques 2007**

Amsterdam, The Netherlands
www.nikhef.nl/acat07/

24–27 **HealthGrid 2007**

Geneva, Switzerland
<http://geneva2007.healthgrid.org/>

May

7–11 **OGF20 and EGEE User Forum**
Manchester, UK
www.ogf.org/
www.eu-egee.org/uf2
OGF20 proposal submission deadline: 9 February
EGEE User Forum abstract submission deadline: 31 January