



CERN COMPUTER NEWSLETTER

Volume 41, issue 2 April–May 2006

Contents

Editorial

AIS group keeps CERN on track **1**
Computing featured in this month's *CERN Courier* **2**

CNL archive

Computing at CERN in the 1970s **3**

Announcements & news

New LDAP server for e-mail address book **4**
CRA account tool comes into effect **4**
CERN urges mail clients to use secure protocols **4**
Oracle support provides a range of new tutorials **4**

Physics computing

New tape storage for LHC era **5**
Linux update: CERN to migrate to SLC4 this year **5**

LCG news

The view from the experiments **6**

Conference reports

CHEP'06 anticipates the start of LHC experiments **8**
EGEE User Forum marks a new phase for Grids **9**


Technical briefs

BOINC activities at CERN **10**
CERN Windows Terminal Services **11**

Information corner

Working with NonAdmin on your PC **13**
CERN School of Computing will be held in Finland **13**
Calendar **13**

AIS group keeps CERN on track



The Earned Value Management system (left) is tracking various aspects of the cost of the Large Hadron Collider (above) through to its completion.

handled 250 000 electronic documents, involving 650 000 digital signatures. While most staff and users at CERN may be familiar with EDH, and many know about CERN Expenditure Tracking [CET] and Human Resource Toolkit [HRT], they often don't realize that there is a core of corporate systems that makes all of this work.

In fact, EDH is responsible for the workflow of documents – routing them to receiving systems. There are some seven such systems, depending on how you count them: one deals with financial accounting and purchasing, another deals with human resources, and so on. If we lost EDH we would in theory lose nothing that is absolutely vital to the organization, since all of the data ends up in these core systems. If on the contrary we

Each time you fill out an electronic document handling (EDH) form on the Web or receive a pay statement electronically, the Administrative Information Services (AIS) group will be involved behind the scenes, deploying, maintaining and often developing the software that will deal with them. CNL spoke to the AIS group leader, Reinoud Martens, about the progress that has been achieved over the past few years to streamline these services, and the challenges that lie ahead.

Can you give us a sense of what EDH at CERN involves, in terms of IT infrastructure?
Last year the EDH system

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.

Advisory board Wolfgang von Rueden (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 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



©2006 CERN

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

lost our financial system, no payments (including salary payments) could be made.


AIS uses a combination of bespoke solutions and off-the-shelf commercial products. What determines your software choices?

We have bespoke developments, such as EDH, CET and HRT, when we don't find a suitable product on the market. But where we can buy off the shelf, with an acceptable level of quality and functionality, we do so. This is the case, for example, with the e-recruitment application, which handled more than 19 000 applications last year. The CERN stores are using BAAN for logistics and warehouse management. This is essentially a small business, with a turnover of over SwFr 40 million per year and more than 14 000 different items in stock, so it makes sense to use a standard business solution.

Qualiac is a commercial product that we use for finance and purchasing, as well as goods reception and all movements of goods. It is used internally on the CERN site, and to and from any location in the world. CERN's requirements in these areas are sufficiently standard that a good commercial product can handle them. There are always a few cases where a bespoke solution might provide more or better functionality, but the extra cost usually outweighs the benefits.

A big change over the last few years has been the introduction of increasingly sophisticated project-management tools. Can you sketch this evolution?

Over five years ago we started developing Project Progress Tracking [PPT] for ATLAS and CMS, as a collaborative tool on the Web. This was limited to technical milestones, and did not include financial data. In 2002, as a result of the Aymar review, there was a big push to improve project management for CERN in general and the Large Hadron Collider [LHC] in particular. We started with a crash implementation of PPT for the LHC, loading PPT with the data on contractual financial milestones from our financial systems on which we lacked updated planning information.



Reinoud Martens, group leader of the IT department's AIS group.

This took about four months and was such a success that it evolved into our Earned Value Management [EVM] system in 2003. The LHC project has relied on this tool ever since. We now have 12 000 work units in EVM, each representing a piece of work with clearly defined deliverables, typically over a maximum of a few months.

PPT is a framework – a set of components – that enables us to quickly implement solutions for other projects. For example, Enabling Grids for E-science and CERN Neutrinos to Gran Sasso both rely on PPT. This software is also at the heart of the Activity Planning Tool [APT], a CERN-wide tool that is not restricted to project management, and that was successfully deployed last year. What APT means in practice is that it will become much easier to establish a link between what departments ask for, what they get, and what they actually use – and all of this from a single data repository.

One of the interesting trends that we are seeing with PPT is that other labs and major projects are expressing an interest in adopting it. Both ITER [the international fusion reactor] and the International Linear Collider are seriously considering PPT for their needs.

What plans do you have for AIS in 2006 and beyond?

A milestone for 2006 is the introduction of the Computing Resource Administration [CRA]. This provides far better control with a direct link to the Human Resources database, and it can

be updated promptly when people's roles change or users leave CERN.

Keeping computing accounts and access up to date, and complying with best practice on this, is particularly challenging at CERN where we have a large turnover of short-term visitors. There are more than 170 000 people registered in CERN's Oracle Human Resources database, which is incredible when you think about it. CRA in its first release is really about integrating existing functionality rather than creating new options. This then constitutes the foundation for new functionality that will follow soon, and that will facilitate the use of many services provided by IT, such as single sign-on and coherent e-groups throughout all IT services.

Another significant project will be the integration of the pension fund. We will in the first instance offer a human resource management and payroll service for CERN pensioners, and at a later stage we may extend to financials for the pension fund.

One of the questions we need to keep asking ourselves is: should we continue to combine best-of-breed products from different vendors, or should we go with a single integrated solution from one vendor? This year we will conduct a business study to see if there is a case for changing our ways. At the moment we are burning around 80% of our resources on maintenance and support and have only 20% to respond to new requests, so anything we can do to reduce maintenance is of interest.

Computing featured in this month's CERN Courier

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

Computing News

● **BIG GRID gets grant for e-science**

The BIG GRID consortium has been awarded €29 million to build an e-science infrastructure in the Netherlands.

● **apeNEXT computers take up residence in La Sapienza**

The apeNEXT supercomputer laboratory in Rome has 8 Tflops power and 1.5 TB memory.

● **Volunteers use own PCs to analyse cosmic dust**

Stardust@home participants are actively involved in the search for interstellar dust.

● **Computing Grid succeeds in one gigabyte-per-second challenge**

The Worldwide LHC Computing Grid collaboration completes LHC service challenge.

● **EGEE goes global**

EGEE's Grid extends to China, Latin America, the Baltic States and the Mediterranean area.



● **Open Science Grid elects managerial team**

The new team will lead the OSG's operation and expansion.

Calendar of events

Feature articles

● **Networking tackles the LHC challenge**

The requirements of the LCG have changed the way that HEP networking is organized.

● **Counting on uncertainty**

Quantum computing: between physics and computer science.

Computing at CERN in the 1970s

Following last issue's article "CNL celebrates 40th birthday", we will continue to delve into the archive to bring you curiosities from the early years of computing at CERN. Last issue we reprinted stories and photographs from the 1960s, and this issue we will turn the spotlight on the 1970s.

Ferranti EP140 microfilm plotter
 "For the first time as far as we know, the EP140 at CERN has been used to make a motion film. A description is given in a paper by P Lapostolle and R Le Bail, entitled 'Two-dimensional computer simulation of high intensity proton beams', which was presented at the First European Conference on Computational Physics, Geneva, 10-14 April 1972." (May 1972)

Who uses the computer time?
 "With the intention of matching some real numbers against people's guesses as to who uses the time on central computers, the attached figure [not reproduced] shows the CP time used by each division of CERN in 1974 up to 1st December... 80% of the used time goes to the physics divisions, NP (46%) and TC (33%)..."

"The CP time attributable to the computing 'divisions' of CERN... amounts to some 11% of which a fair fraction should be considered strictly as systems overhead, and a further fraction is work performed for physics groups. Accelerator divisions have used a combined total of ~7%... The Theory division have used only 1.8%. Clearly the theorists are thinking rather than computing." (December 1974)

Archiving and rationing of tapes
 "At present over 100 000 reels of magnetic tape (60 000 in active storage, 10 000 in closed shop, 15 000 in 170 traffic slots and 17 000 in archive) are stored for users by DD in the 513 computer centre building. All available storage space is presently full, and it is only through the good will and cooperation of a few users, who have large numbers of tapes,



The magnetic-tape storage vault in the Computer Centre (March 1974).



The computing centre was enlarged with the installation of a CDC 7600, which is shown here during its assembly in February 1972.



The remote input/output station in building 112 was used to submit jobs to the CDC 6500 and 6600.



The prototype SPS control desk is shown under test with the NORD-10 computer (December 1973).



A course in computer science organized by the Training and Education Service (January 1975).



The Computer Science Library (CSL) in the Computer Centre (December 1973).

that we have been able to survive over the past year and a half and still allocate tapes and corresponding slot numbers (VIDs) on request... The computer Users Advisory Committee agreed in 1975 that space in our 513 vault should be rationed to an approximate limit of 50 000 slots for active tapes. We are well above this limit and a tape archiving exercise... will be performed in two steps." (May 1976)

For messy programmers
 "The first reaction of someone who is presented with a 'finished' or 'working' program is normally complete amazement that such a frightful lash-up could ever succeed. So why not make all your source code look as if it really was designed that way by making it as legible as possible? A new facility, INDENT, is now available on all CDC machines which automatically indents DO-loops by 5 columns, so that in nested DOs the statements begin in columns 7, 12, 17... depending on the depth of nesting." (April 1978)

Online Computer Newsletter (OCNL)
 "This is a new CERN DD publication which will contain articles and news items of interest to users of the various mini- and midi-computers in use at experiments around CERN. The intention is that it should fulfil the same function for such smaller computers as the CNL does for the large central CERN IBM and CDC computers." (January 1979)

NAG program library now available at CERN
 "The NAG (Numerical Algorithms Group) Program Library has now been installed on the CDC 7600 and on the IBM system at CERN. This library is a large and coherent collection of high-quality subroutines for numerical and statistical analysis. It is the fruit of many years' work by nearly all the British universities, and is widely distributed (commercially) around the world." (May 1979)

New LDAP server for e-mail address book

The Lightweight Directory Access Protocol (LDAP) provides an address-book service for mail clients other than Outlook and Webmail. Following its announcement at the Desktop Forum in November 2005 (<http://indico.cern.ch/conferenceDisplay.py?confId=346>), the new LDAP server has been progressively put into production since 15 February. It replaces the outdated existing service (both its hardware and software).

LDAP uses a redundant design for improved availability. Over the last few months it has been tested intensively by several

end-users, and we expect the transition to be transparent.

The service is based on the LDAP standard schema, which supports different-use cases with LDAP-enabled applications. Although the content is slightly different, it will work as before for searches related to e-mail addresses. Since March the LDAP data has been managed and retrieved from the new Computing Resources Administration (CRA) system.

The service has been tested with Mozilla, Mail for Mac OS X and Pine. Some minor side effects with Pine have been documented (<http://cern.ch/mmmsservices/>

[Help/?kbid=022030](http://cern.ch/mmmsservices/Help/?kbid=022030)).

Please note that LDAP supports authentication using NICE log-in and password. As agreed by the Desktop Forum, from June a successful authentication will be required to access the directory from outside CERN. Anonymous access is still possible from CERN's internal network. Details about how to enable LDAP authentication on your mail client are available from the Web page mentioned above.

Any problem with LDAP should be reported to the Helpdesk or to mail.support@cern.ch.

The CERN mail team

CERN urges mail clients to use secure protocols

E-mail is an important vector of viruses and is the origin of many security incidents that affect the whole user community at CERN.

One preferred way for viruses to propagate is through the use of non-authenticated mail protocols. In addition, non-encrypted protocols send users' credentials over the network in clear text, which may enable passwords to be "sniffed". Therefore we strongly recommend using authenticated and encrypted protocols for sending and retrieving messages, as well as for using the mail address book.

These protocols are already the default for CERN Outlook users. They can easily be activated in Pine, Mozilla and Mac Mail by changing the configuration to enable a security option (SSL or TLS) for each affected protocol: SMTP for sending, POP/IMAP for retrieving, and LDAP for address look-up.

For details about the appropriate configuration, see the help page at <http://cern.ch/mmmsservices/Help/?kbid=191040>. Most users will also find tools there that will enable them to update automatically the settings of their mail client.

Over the coming weeks, if your client is not securely configured you may receive an e-mail from Mail-service@cern.ch with instructions on how to correct it.

From June, access to mail servers from outside CERN will be restricted to mail clients that are correctly configured, as described in the help page mentioned above. The long-term aim is to stop all use of insecure mail protocols.

The CERN mail team

CRA account tool comes into effect

In the November 2005 issue of *CNL* we gave you early warning that new software would be introduced to improve the administration of computer accounts. This software, called Computer Resources Administration (CRA), was introduced in mid-March, replacing the nearly 20-year-old Oracle-based Computer Centre Database (CCDB) and its user interfaces, *userinfo* and *userreg*.

The new interface for users is an Administrative Information Services (AIS) Web application, <http://cra.cern.ch/>. To use the application you will need to authenticate your AIS account. AIS accounts have been created for users who already have a NICE account and use a NICE user name and password. People without an AIS account can ask their group administrator to

create one for them. New users will automatically get an AIS account when they first register.

The options displayed by CRA will depend on the role of the authenticated user. End-users will see at the left of the screen the following options:

- User Info, for displaying your account information (My Details) and the option of updating your e-mail address and nickname;
- Quota, for requesting a change to your AFS quota;
- Service Account Request.

In addition, Computer Group administrators will be able to create accounts, protect special accounts, and manage AFS space.

By the time this issue of *CNL* is published the automatic account clean-up features will probably be enabled. These include the blocking of unused accounts and those of people

whose registration at CERN has expired. The clean-up features will be applied first to the IT Department and later to the rest of CERN. Warning e-mails will be sent to users and their supervisors at least one month before any accounts are blocked. For more information, see www.cerncourier.com/articles/cnl/2/11/6/1.

The CRA and Registration teams

● **Please note:** in order to receive warnings, external users are encouraged to use the CRA interface to check their registered e-mail address and ensure that it is correct. It is believed that some users have the wrong address registered in the central database, which is used by both the Users Office and an increasing number of applications.

Oracle support provides a range of new tutorials

The IT DES Oracle Support team is pleased to announce a new series of Oracle tutorials, with the proposed schedule:

March

30: Design, Arash Khodabandeh.

April

20: SQL I, Eva Dafonte Perez.

27: SQL II, Lucia Moreno Lopez.

May

4: Architecture, Montse Collados.

11: Tuning, Michal Kwiatek.

June

1: PL/SQL I, Eva Dafonte Perez.

8: PL/SQL II, Nilo Segura.

15: Oracle tools and bindings with languages, Eric Grancher and Nilo Segura.

These tutorials will take place in the IT Amphitheatre (building 11/3-004) and will start at 10 a.m. They will last about one hour, with extra time for

questions. There is no need to register in advance.

You can access the 2002–2003 sessions at <http://it-des.web.cern.ch/IT-DES/DIS/oracle/tutorials.html>.

For more information, please contact Catherine.Delamare@cern.ch.

IT DES Oracle Support team

New tape storage for LHC era

The Large Hadron Collider (LHC) will begin operation in 2007 and is expected to produce 15 PB of experimental data each year. At least one copy of the data will be stored on tape because this method is still cheaper than disk.

At the beginning of 2005 CERN's tape-storage infrastructure for experiments was provided by 10 Sun StorageTek Powderhorn silos and 44 9940B Sun StorageTek tape drives. This infrastructure can store around 4 PB of data using approximately 20 000 9940B tapes. The Sun StorageTek Powderhorn is due to reach the end of its life in 2009, and therefore the replacement of the robots and tape drives has to be planned to avoid disruptions during the operation of the LHC.

What are the options?

The process of selecting tape technology for LHC was initiated in the first months of 2005. The aim was to determine which equipment would best suit the computing requirements of CERN's high-energy physics (HEP) environment, which has traditionally involved a large amount of random reading of data from tapes. This differs from the common tape-drive usage pattern of data back-up, where data is written once and rarely read back. It is widely argued that lower-cost drives (such as LTO, which is popular in the world of "back-up") are



The Sun StorageTek SL8500 tape library in the Computer Centre.

not entirely suited to CERN's patterns of use.

Few companies share the high end of the tape automation market. Two tape robots were selected from this restricted group for testing at CERN's computer centre. One is the IBM 3584 tape library, which in its maximum configuration can contain about 6000 cartridges and supports IBM's high-end and LTO tape drives. The second is the Sun StorageTek SL8500 tape library. In its maximum configuration this can contain approximately 300 000 cartridges and accommodate 2048 drives. It supports recent Sun StorageTek drives, as well as LTO and SDLT tape drives from other vendors.

The current candidates for the tape drives that will record LHC experimental data are the enterprise-class drives from IBM and Sun StorageTek. These are the IBM 3592 E05, which has a

native data rate of 100 MB/s; and the Sun StorageTek T10000, which has a native data rate of 120 MB/s. Both of these drives use a 500 GB capacity cartridge. The "commodity" LTO3 tape drive is also being evaluated. This offers a native data rate of 80 MB/s and a 400 GB capacity cartridge.

We now have two test libraries at CERN. One is an IBM 3584 library configured with some 6000 storage slots and equipped with 40 IBM 3592 E05 tape drives. This represents the basic building-block of a potential IBM solution for LHC data storage, which would consist of as many of these independent libraries as was needed. The second is a Sun StorageTek SL8500 that is configured with about 8000 storage slots and equipped with 12 T10000 tape drives and 21 LTO3 tape drives. This represents the basis of a potential Sun StorageTek

solution to the same problem. The configuration exploits a difference between this and the IBM option, as the basic Sun StorageTek libraries can be interconnected by a "pass-through" mechanism that enables the assembly of a much larger single library. To test the pass-through capabilities of the SL8500, some 4500 9940B cartridges and 10 9940B drives have been moved into this library and are running normal production.

These configurations are being used to determine the performance, robustness and ease of administration of robots, tape drives and media so that we can choose the most suitable equipment for storing data from the LHC experiments. We are particularly interested in finding how to best integrate these high-speed storage devices into CERN's CASTOR Hierarchical Storage System to profit from their different characteristics and achieve the optimum storage performance.

The definitive outline of the future configuration of CERN's tape infrastructure will only be known later this year. What is already certain is that tape, which has been used for storage since the introduction of computing at CERN, will continue to play an important role in the storage of LHC data, whichever equipment is selected.

Hugo Caçote and Charles Curran, IT-FIO

Linux update: CERN to migrate to SLC4 this year

In the November–December 2005 issue of *CNL* ("CERN explains its Linux strategy: SLC4 or SLC5?") it was hoped that Scientific Linux 5 (SL5) would be fully certified for use by autumn this year, enabling us to skip the widespread deployment of SLC4.

However, Fermilab and CERN Linux experts now expect Red Hat Enterprise Linux 5 (RHEL5) and hence SL5 to be available only at the beginning of 2007. Hence, the CERN-certified variant of SL5, SLC5, could at the earliest be available at the

end of the first quarter of 2007. This is clearly too late to provide a stable start-up Linux platform for the Large Hadron Collider (LHC), and it is equally clear that it is not possible to stay with SLC3, as this is supported only until the end of October 2007.

SLC4 for LHC

Therefore time must be found in this year's busy schedule for a migration from SLC3 to SLC4. This does not imply that all of the LHC Computing Grid sites also need to migrate to a new

operating system, and still less that they all migrate to SLC4. However, even a CERN-only migration requires co-ordination with our partners.

The CERN Linux team wants to raise awareness of this issue, and discussed it with the LCG Grid Deployment Board (GDB) at its meeting in March, thus giving time for sites and experiments to prepare. The CERN Linux team hopes that the GDB will be able to agree on a time window for migration at the latest at its April meeting, which will be

held jointly with HEPiX.

By the time you read this, SLC4 should be certified (see <http://cern.ch/linux/slc4> for exact status).

For an overview of the Scientific Linux releases and the support periods for each of these releases, see the Distribution Roadmap page at <http://www.scientificlinux.org/distributions/roadmap>. Please note that this is just a plan. It cannot be guaranteed because there are too many factors that may affect it.

The CERN Linux team

The view from the experiments

This article is based on interviews with computing co-ordinators from ALICE, ATLAS, CMS and LHCb. They shared their thoughts on their use of the LHC Computing Grid (LCG) to date, as well as the challenges that lie ahead.

The good news

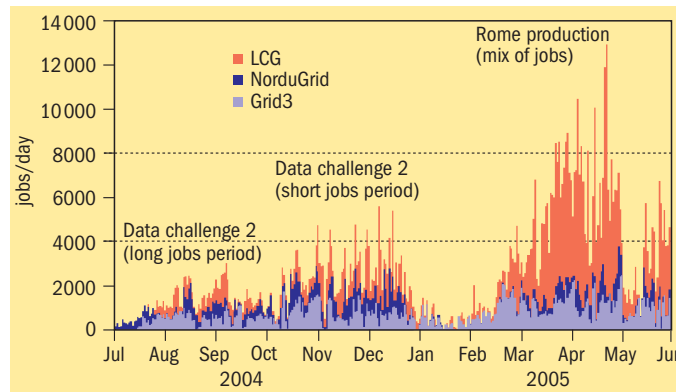
LCG was launched as a service in September 2003, and since then much has been achieved. One of the quantitative highlights has been the production by all experiments of tens of millions of Monte Carlo simulations using LCG, NorduGrid and the Open Science Grid (OSG) – the three types of Grid that underlie computing for the LHC.

Dario Barberis, the computing co-ordinator of ATLAS, explained that ATLAS produced 10 million events in 2004 during Data Challenge 2. From January to May 2005 it produced about 12 million more events, which were analysed at a workshop held by ATLAS in Rome last June.

“One of the main achievements of the Rome production was that we ran with peaks of more than 10 000 jobs a day,” said Barberis. ATLAS is preparing another large production for this spring, when the detector’s real geometry (including the sagging, distortions and misalignments that are being measured as the detector is assembled) will be used to test the alignment procedures and calibration algorithms under realistic conditions.

CMS has used LCG and OSG to produce about 60 million Monte Carlo events, to prepare the *Physics Technical Design Report*, said CMS computing co-ordinator Lothar Bauerdick. The CMS collaboration is now starting to push this data through the Grid to simulate analysis. The analysis is revealing fragilities in the infrastructure, but, said Bauerdick, “people are successfully running jobs, especially if they are not afraid of the technology.”

It is perhaps surprising to discover that LHCb – the smallest of the four collaborations – has so far been the major user of LCG’s CPU resources. LHCb produced in



ATLAS record of the number of computing jobs per day from July 2004 to June 2005, with peaks above 10 000 during the Rome production.

the order of 100 million events in 2005 alone. LHCb began to run Monte Carlo simulations in 2003, accessing resources mainly outside LCG using Dirac, LHCb’s workload management system. “Over time this has switched, and we are now over 95% on LCG,” said Nick Brook, LHCb’s computing co-ordinator. The experiment even plans to use its online farm of 1800 machines, comparable to some LHCb Tier-1 centres, for data reprocessing during shutdowns.

Using an agent-based approach to submitting jobs has been very successful for LHCb. “What we submit to LCG is not the jobs themselves but an agent, which queries the central Dirac system, asks ‘Have you any work for me?’ and downloads and runs the job,” said Brook. The wrapper around the agent tests the local environment to find out, for example, if there is enough disk space, and if the relevant software is installed. If a job dies, only the agent is lost. “This builds an extra layer of robustness into our model,” said Brook.

ALICE has successfully tested its I/O solutions for access to data, and has achieved a good workload distribution by combining LCG with its own central task queue technology. During its data challenges ALICE has achieved 3700 jobs running in parallel. “This is significant because these are substantial jobs, typically taking eight hours or more to complete,” said Federico Carminati, the computing co-ordinator of ALICE.

Carminati recognizes that working closely with the IT Department during data challenges has helped to achieve such targets and increase confidence in LCG. “People have been extremely helpful and responsive,” said Carminati. “This has been a very successful collaboration.” He added that the belated but high-quality release of CASTOR-2 was a critical step forward for ALICE. This experiment must push data rates to mass storage from hundreds of MB/s, the level so far achieved, to 1.25 GB/s, which is anticipated in ion runs.

Lessons learned

Looking at the bigger picture of how LCG has evolved, Carminati says that “we now have a solution that is both intellectually and practically satisfactory.” He views as key the decision by LCG’s management to establish a baseline services working group so that experts could focus on constructing a solid foundation. Rather than trying to provide all imaginable functionality at once, the targets have been file transfer, workload management, accounting, security, monitoring and information systems. These basics have to be rock solid and transparent to users if the experiments are to build on this.

As time goes by, more of the higher-level functionality that the experiments are now developing independently will move to the base services and be shared. Everyone will benefit

from this because redundant functionality will be eliminated, thereby freeing up resources in the experiments. “This is one of the big gains of the Grid,” said Carminati. However, he regrets that it took so long for all involved to focus on the essentials. “Everyone thought they could provide a complete solution easily,” he said, “but the problem was much more complicated than anyone expected.”

For Brook, the issue now is to hide the complexity from end users, so that the switch from traditional modes of analysis to submitting via the Grid occurs rapidly. LHCb and ATLAS are developing a user interface called GANGA, supported by both GridPP and EGEE, to handle this. “Once we convince users to switch to GANGA, the details of the back-end will no longer matter, whether it’s LSF, another batch system or the Grid,” said Brook. He noted that GANGA test users are knowledgeable and have intricate use cases, which tends to complicate matters, “although it is useful for shaking down the system”.

Barberis pointed to another learning curve that he has had to deal with; namely, the different expectations of the Grid middleware developers and the people developing software in the experiments. “Some of the functionalities we had hoped to get two or three years ago are still not in sight because these things were low priority at the outset,” he said. The need for them is becoming much more urgent as the experiments approach full production, but the developers are not necessarily aware of this.

As an example, Barberis cited the functionality of sending a job to where the data is. This was requested during EDG, the DataGrid project that preceded EGEE, but was always considered low priority. This was because data transfer was not critical for the large Monte Carlo simulations that characterized early Grid use; these required relatively little data and great amounts of CPU. But for analysis jobs, which take perhaps a

thousandth of the time of simulations, “you can spend most of your time waiting for a file to arrive”, said Barberis. “We want the job to go where the data is, even if it has to wait there a while, rather than congesting the network with useless data transfers.”


Getting this sort of message across to the middleware developers is mainly a question of improving communication channels. Barberis is optimistic that the recent creation of the EGEE technical co-ordination group, a forum that provides more direct contact between the experiments and the developers, will help with this.

Bauerdick pointed to another key lesson learnt from using LCG: it is important to start tests even when the system is far from perfect. Users at Fermilab have been encouraged to test the CMS system even though job completion rates are low, at around 70%, owing to various issues, some of which are on the CMS side. One unexpected result of such testing occurred when, owing to a bug, users produced 10 times their anticipated I/O. “We eventually fixed the bug,” said Bauerdick, “and in the meantime we stress tested our I/O, and this also helped us debug our storage systems.” While such accidents are hard on today’s users, they are beneficial in the long term.

The benefits of WLCG

An important development remarked on by several computing co-ordinators was the establishment last November of the Worldwide LHC Computing Grid (WLCG) collaboration. The WLCG unites the stakeholders in LCG/EGEE, OSG and NDGF (the Nordic Data Grid Facility) in a collaborative framework that is similar to the one adopted by the experiments.

Barberis noted that ATLAS is a big user of all three infrastructures, and said that each Grid infrastructure catered to a different type of centre. “The NorduGrid software has advantages for small institutes because it is easy to set up. With one machine you can realize a gateway between the Grid and your centre,” he said. In contrast, LCG’s software is more invasive but better



A simulation of a Higgs event produced during a CMS data challenge.

adapted to being deployed on dedicated large facilities, because it provides a more uniform environment. “OSG has much in common with LCG/EGEE, but OSG decided not to develop a higher-level layer,” he said. This means that there is no resource broker – ATLAS users simply decide which sites to use when they submit a job.

Bauerdick is pleased with the progress in interoperability that WLCG encourages. “LCG has really embraced interoperability,” he said, “and as an experiment we are already using both LCG and OSG for analysis.” However, he argued that although the collaboration is working well, “it is important that participants have the spirit of doing the LHC, not just doing the Grid”. Bauerdick is trying to bring operations staff in Tier-1 centres to CERN so that they can see the detectors, “to show them how wonderful this is, so they really buy into getting the LHC computing working”.

Immediate challenges

The computing co-ordinators tend to agree on the two big challenges facing LHC computing over the next 18 months or so, in the run-up to first beam.

The first challenge is to provide a stable Grid service. According to Barberis, “robustness is difficult to achieve in a Grid environment, and we still suffer from a mishap somewhere in the world bringing the whole system down.” Brook is worried that “centres that have never run 24/7 must now provide a stable service, which means they cannot take down services at minimal notice, they have to

stick to the service agreements”.

The second challenge is to scale up from today’s small numbers of test users – typically 20 or 30 for each experiment – to the large numbers expected once data starts to flow. For LHCb, more than 150 people could become active on the Grid; for ALICE, it is several hundred; and for CMS and ATLAS the number is probably about 1000. Part of the problem is that there is a gulf between the technophile test users who are debugging the systems and the professors in remote institutes who are trying to run jobs for the first time.

“We do not know whether the typical end user will have efficient access to the data; whether the analysis will be easy or difficult,” said Carminati. ALICE is just starting to open up for users, with up to 100 expected over the coming months. “We really need to understand the pathologies,” said Carminati. For example, what happens when a non-expert user accidentally submits an inordinate number of jobs?

Bauerdick is concerned that with the many different layers of software involved, it will be difficult for new users to know who to talk to when problems occur. “We need an end-to-end triage system to deal with queries,” he said, “and we lack the ability to do error propagation to the point where you can solve problems.”

Carminati highlighted another potential problem: “experience shows that every time we reach a new threshold for the number of users of a system, we discover new limits. Most of the

time we can fix these, but what if we discover a fundamental limit that invalidates our design? We cannot fix this by just adding horsepower.”

Countdown to physics

Despite these concerns, the sense of excitement and urgency is palpable in the computing co-ordinators’ descriptions of their plans for the months ahead. For example, both ATLAS and CMS have already begun measuring cosmic rays. CMS will start to see muons going through several detectors by April. For ATLAS, a major cosmic ray run is planned for the autumn, when more detectors are installed. The data will be at a rate of about 1 Hz rather than the 200 Hz planned for when the LHC starts up, but, as Barberis said, “it is critical that the offline computing is fully operational for this, since this is real ATLAS data.”

Looking beyond this, Bauerdick has his sights on the start-up of the LHC. “Even in the first few weeks we will be looking for exotic signatures,” he said. “The difficulty will be to understand the detector well enough to know whether what we are seeing is real.” Although the Higgs particle is unlikely to turn up in the early data, the supersymmetric squark and gluino, should they exist, are expected to give rise to spectacular signatures in the detectors.

CMS and ATLAS are not alone in expecting early results. Brook said that one of the things LHCb will be looking for is $B_s - \bar{B}_s$ oscillations with periods up to about 70 ps. “If the Tevatron does not beat us to this, this may be an indication of new physics beyond the Standard Model,” said Brook.

ALICE will also be looking for interesting events from day one. As Carminati noted, “we have an excellent tracking system, which means that we are in a good position to exploit the first low-luminosity proton runs.” These runs will also help to prepare for the much more complex ion runs, which could start at the end of 2007.

Bauerdick summed up the mood when he said, “we cannot afford to use the first data to test the computing models, because that data may be the key to a discovery.”

François Grey, IT/DI

CHEP'06 anticipates the start of LHC experiments

Computing in High Energy and Nuclear Physics (CHEP) is a major series of international conferences for physicists and computing professionals in the fields of high-energy and nuclear physics, computer science and information technology.

CHEP'06 took place on 13–17 February in Mumbai, India. It was held in the Tata Institute for Fundamental Research (TIFR), which is celebrating its 60th anniversary this year. An unusual event at the end of the conference was a visit by the Indian president, H E Dr A P J Abdul Kalam.


The organizers estimated that about 500 people attended the conference. This included several who only attended the pre-conference LHC Computing Grid (LCG) Service Challenge meeting, and others who participated in the CHEP'06 preconference school. As was to be expected, more than 50% of the papers presented at the conference referred to preparatory work for one or the other of the Large Hadron Collider (LHC) experiments.

Opening speeches

The conference was opened by the TIFR director, Prof. S Bhattacharya, who spoke of the institute's history and how and when it became involved with high-energy physics (HEP), among other areas of science. He noted that this would be the last CHEP before the start-up of the LHC, and said that TIFR was proud to host the conference.

The TIFR director was followed by several other local dignitaries who noted the importance of grids and Grid software for modern experimental teams – in HEP and in other disciplines, and increasingly beyond science. It was noted that TIFR is a Tier-2 site for CMS, with a gateway for the ALICE group in Calcutta.

Opening the plenary sessions, CERN Chief Scientific Officer Jos



The president of India, Dr A P J Abdul Kalam, gave the farewell address.



A delegation from CERN visited the president in Delhi after CHEP'06.

Engelen spoke about the status of the LHC. First of all he detailed the progress made in installation and testing, and the pilot run that is planned for 2007, and then he talked about the status of the experiments.

The following are highlights of the plenary sessions:

- Jamie Shiers (CERN/IT) reported on the readiness of the computing infrastructure for LHC. He said that the technical design requirements of the experiments were now exercised in a series of LCG service challenges in a production-like environment. The most recent service challenge achieved an aggregated 1 GB/s to the Tier-1 sites and up to 200–250 MB/s to

individual sites. Jamie said that he was “cautiously optimistic”, although there was still much work to be done.

- Paris Sphicas (CERN/PH) from CMS described the readiness of the experiments' software, and said that the common software applications area projects (SPI, ROOT, POOL and the simulation tools) were advanced and being used across the experiments. He reported that analysis software was progressing with a common understanding of the format of the data to be used; Geant4 was deployed across all experiments with good progress in data generation; so-called fast simulation was less well advanced and was at different

levels in different experiments; and software performance was becoming the main challenge. Paris confirmed Jamie's fears of a data-intensive start-up despite expected initial low luminosity. He explained the physics that is expected, and noted that at least one experiment (ALICE) is expecting to publish two papers within one to two weeks of the first collisions.

- Beat Jost (CERN/PH) reviewed future data-acquisition systems, both for the LCG experiments and for other, non-CERN planned or projected experiments up to the International Linear Collider in 2017.

- Liz Sexton-Kennedy (Fermilab/CMS) gave a review of event-processing frameworks.

- Martin Purschke (Brookhaven National Laboratory) presented some experiences from the PHENIX experiment at BNL, which is taking data at rates approaching those of the LHC experiments.

- Harvey Newman (Caltech) reviewed the status of worldwide networking for HEP, and noted the recent rapid growth in traffic following the roadmap predicted by various reviewers at previous CHEPs.

- Les Robertson (CERN/IT) reviewed the status of the LCG service, showing the links between LCG and the different phases of the EGEE (Enabling Grids for E-sciencE) project. He compared the performance of LCG to the Gartner curve of highs and lows of performance of a new project over time.

- Ruth Pordes (Fermilab) described Grid interoperability as seen from the US. Interoperability inevitably means collaboration and includes human and social factors, she said. These lead to gateways, resource interfaces and overlapping services, all of which use a common network fabric. This involves in-depth testing, demonstrating that it

works in practice, as well as continual communication between the partners. LCG is a pioneer in showing interoperability between grids on different continents.

- Rene Brun talked about “ROOT in the era of multi-core CPUs”. Assuming that vendors’ claims for multi-cores are correct, there will be an impact on ROOT, he said. However, to use this, applications must be multi-threaded (a future development of PROOF for example). The powerful multi-core PC, or a local cluster of these, is inherently less complex than a grid and could therefore be an easier alternative.

- Wolfgang von Rüden (the head of CERN IT Department) gave a public lecture “From WWW to the Grid” in which he drew parallels on how the latter may one day affect our everyday lives just as the Web does now.


- Randy Sobie (University of Victoria, Canada), the new chair of the International High Energy Physics Computing Coordination Committee, gave a short review of its activities.

- L Kathragadda (from the Indian Google research group) talked about data mining. He noted that power consumption is a major issue and will become more so.

Business talk

Unusually, the CHEP plenary sessions included three “commercial” presentations:

- Tony Hey from Microsoft gave a talk on “e-Science and Cyberinfrastructure” and spoke about Microsoft’s commitment to working with the open-source



The president spoke about collaborations between CERN and India.



Attendees participated in a discussion after the president’s speech.

common standards community.

- David Axmark, the co-founder of the company behind MySQL, described storage engines that have been developed and optimized for different environments, ranging from HEP to banking. He also presented a list of new features that appear in version 5, and other features that will come in future releases.

- Alan Gara of IBM spoke about supercomputing and showed several of the scientific challenges for which high-performance computing is required. He described some recent supercomputers, such as ASCI Purple and Blue Gene (which has 32 TB of RAM).

As noted earlier, the president of India came on site on the last

day of the conference to give the valedictory address. His talk, which was illustrated with screen shots of LCG monitoring, was about the knowledge grid and how it could be used to help in India’s development. The president finished his talk with these words: “My best wishes to all the participants of this conference for success in their mission of removing the digital divide among the scientists located in different parts of the world leading to path-breaking discovery in the area of high-energy and nuclear physics.”

You can browse or print the president’s full presentation from his website at www.presidentofindia.nic.in/scripts/sllatest1.jsp?id=734.

We do not have the space to summarize the many talks that were given in parallel sessions (eight parallel streams were held over four and a half afternoons). The talks covered many areas such as distributed event production and processing, software components and libraries, computing facilities and networking, Grid middleware and e-Infrastructure operation, and online systems.

For details of these talks you may refer to the *Report on the Mumbai CHEP’06 Conference*, written by Alan Silverman and with compiled contributions by M Kwiatek, G McCance, I Papadopoulos and B Tomlin. The report is available on the CERN Computing Seminars website at http://cern.ch/cseminar/2006/0303/trip_report.pdf.

Alan Silverman, IT/DI; Nicole Crémel, IT/UDS

EGEE User Forum marks a new phase for Grids

The first User Forum of the Enabling Grids for E-science (EGEE) project was held on 1–3 March at CERN in Geneva, Switzerland. The event coincided with a significant milestone for the EGEE project, as the total processing power achieved by this Grid infrastructure now exceeds 20,000 CPUs on average.

The forum brought together more than 240 people – both from the established EGEE user community as well as new users from academia and industry – to share their experiences and set new targets for the future.

Discussions included the evolution of existing scientific applications that have been “gridified” to work on this infrastructure. Also discussed was the development and deployment of new Grid middleware and applications in fields such as life sciences, computational chemistry, earth observations, astroparticle and particle physics, and fusion.

The forum was the first in a series of planned events for users that will take place during the second two-year phase of the EGEE project (EGEE-II).

Negotiations with the EU for EGEE-II are very advanced, after the project passed its second review by EU-appointed Grid experts in December 2005. The second phase of EGEE started on 1 April to ensure a seamless service to current EGEE users.

“This event marks the beginning of a new phase in the evolution of scientific Grids,” said IBM Grid strategist Jean-Pierre Prost (IBM Products & Solutions Support Center, Montpellier, France). “The underlying infrastructure is now in place, and increasingly real users

– both academic and industrial – will be driving the further development of the technology, with their expanding range of requirements for scientific computing and data processing using Grid technology.”

The next major EGEE event will be the EGEE’06 conference on 25–29 September in Geneva, Switzerland. This will highlight the range of related projects that EGEE has helped to generate, and the broad spectrum of user communities that are involved.

Hannelore Hämmerle, IT/DI

BOINC activities at CERN

BOINC (Berkeley Open Infrastructure for Network Computing) [1] is an open-source software platform for distributed computing that employs volunteered computer resources. It uses idle CPU cycles of participating computers to perform scientific calculations.

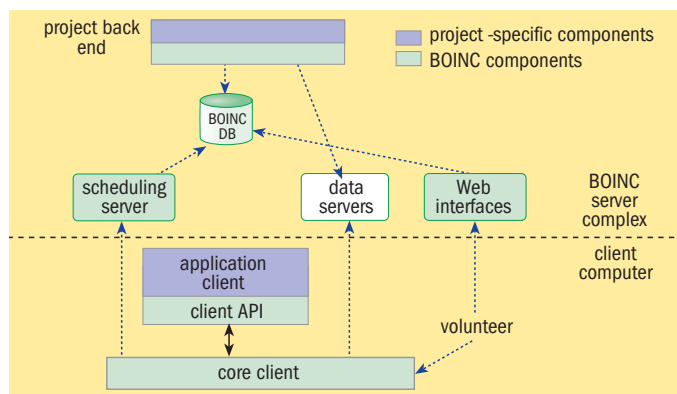
BOINC can be seen as a specialized Grid that is suitable for running applications that are “pleasantly parallel”, so that calculations can be distributed to thousands of machines that do not communicate with each other. Such a configuration is ideal for compute-intensive tasks with modest input/output needs. Nevertheless, the BOINC infrastructure is simple and requires only one server, although this can be distributed to several machines for performance reasons.

The applications that are run in a BOINC project must have public appeal so that participants are willing to volunteer their resources. Outreach is an important part of this kind of project, to give the public a sense of direct participation in a scientific project. Attractive screen savers are often provided, together with a credit-based ranking system for users.

Successful projects

The first project to run on BOINC was the famous SETI@home, which analyses data taken from radio telescopes looking for extraterrestrial intelligence. The project has delivered more than 9 million years of aggregate computing time and has attracted more than 5 million volunteered CPUs. BOINC is now being used by many projects in domains such as physics, medicine and climate prediction.

One of the leading BOINC projects is CERN’s LHC@home [2,3]. Its first application has been SixTrack, which simulates particles circulating around the Large Hadron Collider (LHC) ring to study the long-term stability of the particle orbits. Other applications of interest to high-energy physics are being prepared. At the moment, the



The components involved in a BOINC distributed computing project.

system has about 14 000 active users, 25 000 active hosts and can provide more than 5 Tflops sustained processing rate. This has enabled the computation of more than 700 CPU years in a few months, assuming as typical a 1 K5fp2K CPU that is approximately equivalent to a 2.8 GHz Intel Xeon processor. Of course, this power is not totally free, but the cost:performance ratio is positive, and this is without taking into account the outreach effects. LHC@home still has the potential to become much larger if enough work is available to use the computing capacity that is offered.

In this context a study has been conducted into combining public-resource and Grid computing [4]. Bridges for sending jobs between Grids and public-resource computing projects have been tested and built for LHC Computing Grid [5] and NorduGrid/ARC [6] middleware.

The CoLinux environment that enables Linux to be run as a Windows kernel extension has also been tested for its possible integration with BOINC, to allow jobs to migrate with their own operating system.

One hundred old farm PCs in the Computing Centre that were to be retired have been recycled to run as BOINC clients. We are managing these machines using Quattor with a BOINC client RPM (Redhat Package Manager) that we developed. As well as contributing to LHC@home, the PCs are also enabling studies of metascheduling mechanisms that should make it possible to

implement applications with volunteer resources complemented by limited dedicated ones to guarantee quality of service.

A BOINC server RPM has been produced to automate the installation and management of BOINC servers. We are testing it now and have found and resolved some bugs. The RPM and the bug corrections are being fed back to the BOINC project.

Porting other applications

As we are exploring other applications that may be suitable for BOINC, we have ported the Geant4 [7] release test code to BOINC. This application represents a simplified test beam that is used to validate new versions of Geant4. The port has been done for Linux and Windows in collaboration with Geant4 developers. It was presented at the Geant4 workshop 2005 and at CHEP 2006. The next step is to consider more realistic and useful Geant4 models, and we have already made first steps to achieve that.

We have also successfully ported ATFAST 0.02.22 [8] to the BOINC environment and demonstrated it with a real physics case. This is the FORTRAN version of the ATLAS Fast simulation and reconstruction software, which uses the PYTHIA 6.2 [9] event generator.

The Garfield/Magboltz [10] program for simulating gaseous detectors has been ported to the BOINC framework in co-ordination with the software’s

author. We have demonstrated real-use cases in Linux using BOINC in recycled farm PCs in the Computing Centre. Although we still have to do a Windows port, we expect to receive requests soon from the LHC experiments for a large-scale production campaign.

Preliminary contacts have been held with research groups in France and Japan regarding a future project called Feynman@home [11]. The groups wish to simulate quantum processes relevant to LHC (and future accelerators such as the International Linear Collider) by calculating their Feynman diagrams using BOINC. This is a computer-intensive task that is closely related to event generators.

We have also helped to start up Africa@home with the Swiss Tropical Institute, the University of Geneva and two NGOs. This is in beta tests and will run agent-based epidemiological models of malaria.

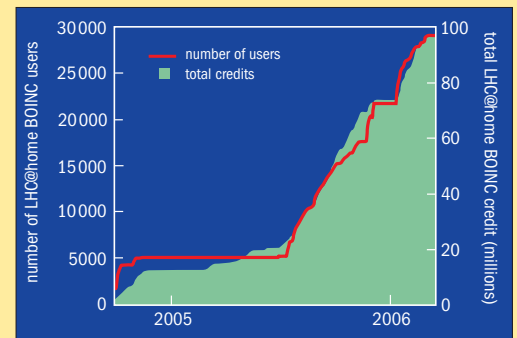
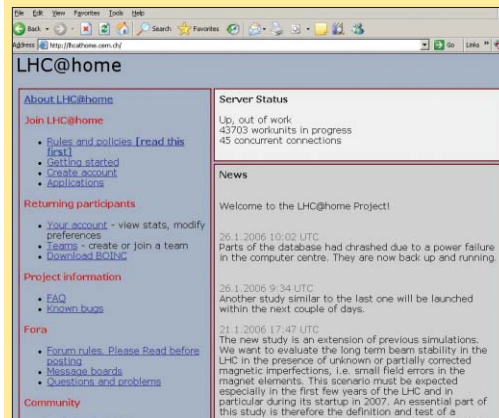
We are also collaborating with the Extremadura region in Spain. The local authorities are going to install BOINC clients in thousands of machines, mostly in schools and hospitals, to run volunteer applications initially provided by CERN. They will also send staff to work at CERN on BOINC infrastructure and application porting. To kick-start the project, a seminar was held at CIEMAT, the Research Centre for Energy, Environment and Technology, in Madrid, attended by representatives from the region. A shorter version of the seminar has been taught as a CERN computing seminar. Slides from these events are included in the LHC@home Twiki area, where we document all BOINC-related work at CERN [12].

We are co-operating with the IS group to deploy the BOINC client to CERN Windows desktops using NICE mechanisms. We plan for a progressive deployment that will start on a voluntary basis in the IT Department. We will subsequently ask for the co-operation of all CERN desktop users in the coming months.


References

- 1 <http://boinc.berkeley.edu>
- 2 <http://lhathome.cern.ch>
- 3 <http://athome.web.cern.ch/athome/LHCathome/whatis.html>
- 4 <https://twiki.cern.ch/twiki/pub/LHCAtHome/LinksAndDocs/ChristianSoettrupBOINCThesis.pdf>
- 5 <http://lcg.web.cern.ch/LCG>
- 6 www.nordugrid.org
- 7 <http://geant4.web.cern.ch/geant4>
- 8 <http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/HIGGS/Atlfast.html>
- 9 www.thep.lu.se/~torbjorn/Pythia.html
- 10 <http://garfield.web.cern.ch/garfield>
- 11 <http://acpp.in2p3.fr/cgi-bin/twiki/bin/view/Feynman/WebHome>
- 12 <https://twiki.cern.ch/twiki/bin/view/LHCAtHome/WebHome>

The CERN BOINC team



The LHC@home website (left), the history of the project's users and total credits (above), and an example of SixTrack output (below).



CERN Windows Terminal Services

What is CERN Windows Terminal Service (WTS)?

CERN WTS enables a remote Windows session to be run from a computer inside or outside CERN's domain that has an operating system that might not be Windows.

WTS started in April 2004. Since then 1300 users have subscribed to it, and around 100 sessions run everyday. There are many reasons for using WTS:

- You may be on an official or private trip and need to access the CERN environment (your calendar, mail or documentation in DFS, for example).
- You want a secure way to connect to CERN.
- You need to run applications that are tricky to install or that consume lots of resources on your local computer.
- You normally work on Linux or Mac but you need occasional access to some applications that run on Windows.

WTS offers easy and secure connectivity, reliability, platform independence and mobility, and helps many CERN users in their work. The WTS website (<http://cern.ch/wts>) describes the service and is the starting point for newcomers.

How can I connect?

Two simple steps need to be taken before using the service:

- You need to subscribe to the service. You can do this on the WTS website. On the Tools panel click on "Register Here" to use the service (figure 1). This will enable your NICE account to connect to the application servers (see "How does it work?" below).

- You need a remote desktop client. Clients for Windows, Mac and Linux are available on the WTS website. For Windows 2000 or XP you can get the latest version of the remote desktop client at the computer's control panel. Select Add/Remove Programs and then look for: Microsoft Remote Desktop connection 5.2.3790.1830 English. Windows XP comes with a client embedded in it. You can launch it by clicking Start / Programs / Accessories / Communications / Remote Desktop Connection. For Linux users, the remote desktop client rdesktop is already included on the CERN Scientific Linux release. Mac users need to download the client from the WTS site. If you have Internet Explorer 5 or higher you can use

the remote desktop Web client available on the WTS site (figure 1) to connect to the service.

A terminal-service-roaming-profile will be created automatically and it will be ready for use. The profile will keep your settings independently of the application server that is holding your session.

To log on to the system we recommend that you use the alias "CERN.TS.cern.ch" (figure 2). The system will then place your session on the machine that can offer the most resources at the time you connect.

If you are connecting from the Internet and have installed Windows Remote Desktop Client 5.2, which is the latest version, you can use TLS 1.0 to encrypt and authenticate the application server (<http://cern.ch/terminalservices/Help/?kbid=020008>). This avoids man-in-the-middle attacks – in which an attacker can read, insert and modify at will messages between two parties without either party knowing that the link between them has been compromised.

Once you have opened a WTS session, your Desktop, Favourites and My Documents

folders are redirected to your home directory.

The Remote Desktop client also enables the use of local resources, such as printers, during the WTS session. For security reasons, the redirection of local drives has been disabled.

If you would like to use a CERN printer that is not already installed, you can request its installation by sending an e-mail to WTS support (support-terminalservices@cern.ch).

What can I run?

The following applications are available during a WTS session:

- Microsoft Office Professional Edition 2003 (many of its components such as Outlook are automatically configured the first time you connect);
- Microsoft Office FrontPage 2003 tool for Web authoring;
- Putty, SSH for Windows;
- Adobe Reader and Acrobat Professional for reading and editing PDF files;
- CERN Phone Book;
- GSView, a post script viewer;
- HummingBird Exceed, for opening an X Windows session on the Windows operating system;
- MS Project Client, a tool for managing resources;

Technical briefs

- CERN Printing Package;
- DreamweaverMX, for Web authoring;
- Whip! Autodesk DWF Viewer;
- WinZip tool, for compressing files;
- HPGL Viewer;
- Remedy Client tool for problem management;
- Microsoft Visio, which creates diagrams and workflows;
- JOracle Initiator, an Oracle plug-in for Internet Explorer;
- Adobe FrameMaker, for creating complex documents that can be distributed to audiences worldwide;
- Adobe Illustrator;
- Nice Alerter;
- WinSCP, for connecting to AFS from the WTS session;
- Java Run Time Environment.

You can request more applications at support-terminalservices@cern.ch. All applications are submitted to the technical criteria of the service. If the application cannot run or is considered unstable on a WTS environment, it will not be installed. Reliability and stability are a must for WTS.

For special requests a cloning service is offered (see our Terminal Service manifest at <http://cern.ch/terminalservices/Help/?kbid=010000>). At the moment there are six cloned application servers.

How long does a session last?

Sessions last as long as they are active. A session will continue on the server when

the client disconnects (either because of a network glitch or because you killed the client). You can always reconnect to the application server when your session is in a disconnected state and continue your work. Please bear in mind that disconnected sessions are only kept for 18 hours. Any unsaved work is lost when the disconnected session ends.

The recommended way to end a session is to log off from the application server in the same way that you would log off on your NICE computer (figure 3).

How does it work?

A light client that implements the Remote Desktop protocol enables a normal desktop computer running Windows, Mac or Linux to connect to WTS and get a remote Windows session. The session is identical to the one a NICE user gets using a Windows computer (figure 4). This means that users can access CERN resources such as DFS, user's home directory and other applications.

The applications you run in the WTS are executed on the Windows terminal servers, and only the display is exported to your computer. That is why you can use this service from any computer that is connected to the Internet, whatever its operating system.

The server infrastructure has four main components:

- Application servers: these are

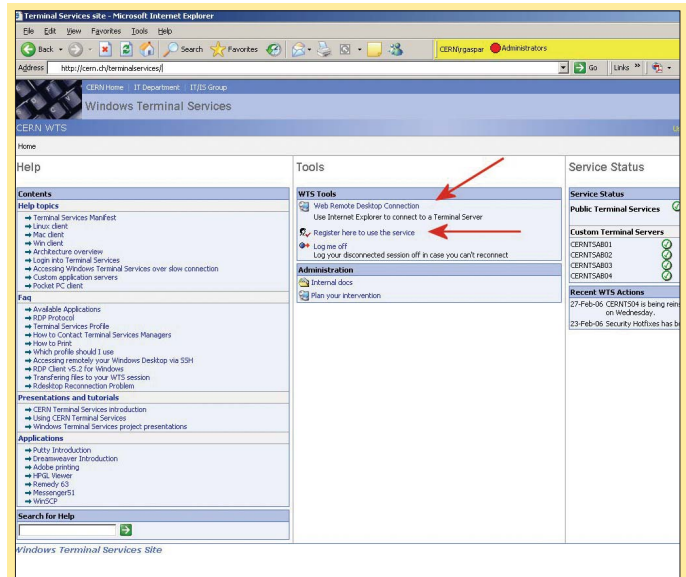


Fig. 1 (above): how can I connect to the Windows Terminal Service?



Fig. 2 (left): when logging in, use the alias CERNTS.cern.ch.

the real machines where users run their session. It is possible to have only one active session on the pool of servers. The communication between your machine and the terminal server is always encrypted.

- Licence server: this enables the transparent connectivity between an application server and a desktop machine running Windows, Linux or Macintosh.
- Directory service: this permits reconnection to a previous disconnected session.

- Terminal-service-roaming-profile server: this gives the impression of a unique environment independently of where you connect. Your terminal service settings will move from one machine to another, making it appear that you are always connected to the same machine.

More information

Please check our website: <http://cern.ch/wts>.

Ruben Gaspar Aparicio, IT/IS




Fig. 3: the recommended way to end a session is to log off from the application server as you would log off from your NICE computer.

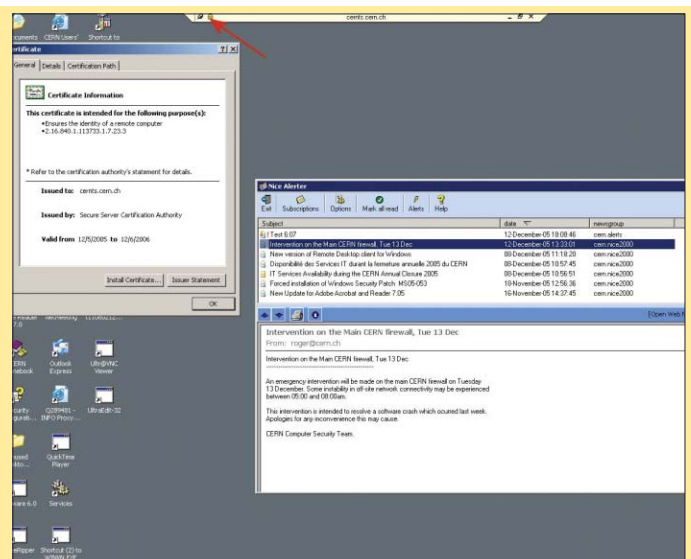


Fig. 4: a WTS session is identical to the one a NICE user gets using a Windows computer. Users can access various CERN applications.

Working with NonAdmin on your PC

As announced in the last issue of *CNL* ("NonAdmin tackles Windows security", p3), Windows XP NICE installations no longer grant administrative privileges by default to the main user or to the person responsible for a computer.

Users without administrative privileges can do most of their work without noticing the change. Common applications (such as Microsoft Office, Internet Explorer, Outlook and Acrobat Reader) will run correctly, and it is still possible to change user-specific settings through the control panel (such as display and regional settings, and Internet options); to configure access to a new network printer at CERN; and to create a new VPN connection to CERN.

There are known issues, however, when you are not your computer's administrator. A work-around is provided for each of these issues, and it is usually based on the NiceAdmin application.

In the Start / All Programs menu on all CERN NICE PCs there is now a NiceAdmin group that has several entries. These are shortcuts to the NiceAdminWin application and options. The first time you launch the application you will need to provide your NICE credentials (only the main

user or the person responsible for the computer is allowed to execute tasks with administrative privileges).

Here is the present list of entries in the NiceAdmin group:

- Add Remove Programs as Admin. This opens Control Panel / Add Remove Programs with administrative privileges.
- Console as Admin. This opens a console window that runs with administrative privileges. Any program you execute from here will have administrative privilege.
- Control Panel as Admin. This opens the control panel with administrative privileges and allows you to choose which setting to change.
- Create NiceAdmin Shortcuts. When using this wizard you can easily create or modify shortcuts to applications so that you can launch them with administrative privileges. The easiest way to use it is to drag/drop the shortcut to the top of your application's list in the Start button (near the Internet Explorer shortcut), then to choose another shortcut (like Remedy), and drag/drop it onto Create NiceAdmin Shortcuts. This will create a modified version of the shortcut so that NiceAdminWin will be called first, giving you administrative


privileges. This shortcut will be placed in the NiceAdmin group.

- Internet Explorer as Admin. This launches Internet Explorer with administrative privileges.
 - Network Connections as Admin. This opens Control Panel / Network Connections with administrative privileges.
 - NiceAdmin Options. This enables NICEAdmin options (such as password caching) to be changed. By default your password is cached in the memory (protected and encrypted) and you have to type it in only once per session.
 - Power Options as Admin. This enables you to change the power options of your computer.
 - Set Time-Date Privilege. This enables you to set the time on your computer using the applet to the right of the taskbar.
 - Account Status. This shows the status of your session (whether or not it is running with administrative privileges) and of your account (whether or not you are a member of the local administrators group). It allows you to permanently change status (for example, to regain or relinquish administrative privileges for your account, and keep the setting across log-ons).
- For more information, visit <http://cern.ch/WinServices/docs/nonadmin>.

CERN School of Computing will be held in Finland

The 29th CERN School of Computing (CSC) will be held at Kumpula Campus at the University of Helsinki, Finland, from 21 August to 1 September, in collaboration with the Helsinki Institute of Physics.

The school is open to postgraduate students and research workers with several years' experience in elementary particle physics, computing or



related fields. The number of participants will be limited to 80. The school is open to applicants from both member

states and non-member states of CERN. Applicants should post a letter of reference and a certification form to arrive by 1 May, and submit a Web-based application form by 15 May.

You can find more details about the school (including programme highlights, practical information, EU grants and how to apply) on the CSC website at <http://cern.ch/CSC>.

Calendar

April

3–7 **HEPIX meeting**
CASPUR, Rome, Italy
www.hepix.org

25–29 **20th IEEE International Parallel & Distributed Processing Symposium**
Rhodes Island, Greece
www.ipdps.org

May

3–6 **Grid and Pervasive Computing (GPC 2006)**
TungHai University, Taiwan
<http://hpc.csie.thu.edu.tw/gpc2006>

4–5 **DEISA Symposium**
Bologna, Italy
www.deisa.org/symposium/index.php

9–12 **GGF17**
Tokyo, Japan
www.ggf.org

15–18 **TERENA Networking Conference 2006**
Catania, Italy
www.terena.nl/conferences/tnc2006

16–19 **Grid Asia 2006**
Singapore
www.ngp.org.sg/gridasia/2006

28–31 **International Conference on Computational Science (ICCS 2006)**
Reading, UK
www.iccs-meeting.org/iccs2006/index.html

30–1 June **First International Conference on Scalable Information Systems (INFOSCALE 2006)**
Hong Kong
www.infoscale.org

September

22–24 **6th WSEAS International Conference on Simulation, Modelling and Optimization**
Lisbon, Portugal
www.worldses.org/conferences/2006/lisbon/smo
Paper deadline: 15 May

25–29 **EGEE'06**
Geneva, Switzerland
<http://cern.ch/eg ee-intranet/conferences/EGEE06>

The deadline for submissions to the next issue of CNL is 5 May

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>