



CERN COMPUTER NEWSLETTER

Volume 41, issue 5 November–December 2006

Contents

Editorial

- UDS provides help to users 1
- Computing featured in this month's *CERN Courier* 2

CNL archive

- Computing at CERN from the 1990s to the Grid era 3

Announcements & news

- CERN mail service requires secure protocols 4
- Printing of banner pages to cease 4
- Discovery Monday profiles the Grid 4
- E-mail addresses at CERN 5

Technical brief

- Virtual servers cut hosting costs and offer flexibility 6

LCG news

- Getting Grids to work together: interoperation is key to sharing 8
- CERN openlab student reports on LCG experience 9

Conference reports

- EGEE'06 debates future of Grid 10
- Digital library solutions employ Grid technology 11
- Openlab event promotes entrepreneurship 11
- CERN is host to EuroPython 2006 11

Desktop computing

- Printing at CERN: procedures for each operating system 12

Information corner

- Let Kerberos guard your password 15
- Known issues for Pine mail client users on Linux 15
- Users can gain access to AIS with NICE password Calendar 16

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.

UDS provides help to users



The human face of the CERN helpdesk forms a small part of the desktop support team, managed by the UDS group and staffed by the firm Serco.

The User and Document Services (UDS) group in the IT department provides a range of support to CERN users, from managing the helpdesk to enabling open access to scientific publications. *CNL* spoke to Tim Smith, the leader of UDS group, about its activities and future plans.

Can you outline the scope of the services that UDS provides?

We are a multifaceted group that covers many domains. The driving projects on the document services side range from the CERN Document Server (CDS), where there are now more than 900 000 bibliographic records and full-text documents, to the copyediting service for CERN reports. We also support paper production activities such as producing CERN reports, running the printshop, and maintaining printers, copiers, faxes and scanners on site. To give you an idea of the scale, the printshop produces some 20 million pages of documents each year, printers produce another 30 million pages and copiers about 10 million.

Then there is the conferencing side of things: we organize meetings and attach documents and minutes to these electronically with the product Indico. We also support CERN's presence in the Virtual Rooms Videoconferencing System (VRVS) and contribute to the worldwide VRVS infrastructure, as well as manage most of the videoconference rooms on site. We organized some 7711 Indico-based events and 3220 VRVS sessions during the first nine months of this year.

We also run all of the auditoria, which entails operator support, maintaining the equipment, taking audio transcriptions of committee meetings, and recording seminars and conferences. UDS also films major events at CERN, which is an important part of the organization's heritage.

When it comes to user support, the main activities are user registration, and management of the central helpdesk and site-wide desktop support through an external contractor. Overseeing the contract is an

active role for CERN; it includes supplying the general manager on duty, and monitoring and training helpdesk staff to ensure that replies are consistent and reliable. The helpdesk receives about 1200 calls each week, so this is a significant task.

Added to that the group still manages the bookshop, which is now located in the main library in building 52. And, of course, we also provide editorial support for *CNL*.

You mentioned user registration – how do you see that evolving?

One of our visions is to automate the process of user registration so that most of it moves upstream – in other words, users can prepare most of the registration process before they come to CERN. This will make the organization appear much more efficient and professional. We would also like to eliminate the use of paper. This would mean ensuring that the process is based on an electronic signature that is integrated into other parts of registration with Human Resources.

What is the status of the scientific publishing service?

Since CERN's inception we have published the well known CERN Yellow Report series to disseminate technical reports and the proceedings of our schools and workshops. At present the in-house copyediting team is not big enough to provide the general publishing service of the past, therefore it is concentrating on continuing the production of the Yellow Reports, which are popular and widely supported.

We have seen a push to rationalize the use of fax machines at CERN recently. What about printers?

For a long time we have had a loosely managed park of printers at CERN, with many different actors buying, operating and maintaining the devices. We want to encourage a more proactive approach, where we can reduce the variety, keep things more up-to-date, guarantee a quality of service, and profit from developments in the outside world in terms of multifunctional machines that can print, copy, scan to e-mail and so on. We aim to develop a proposal for

this over the next few months and implement it next year.

What are the plans for Indico?

We want to make Indico the hub for managing all meetings, which means that all subsidiary systems – videoconferencing, audio conferencing, room booking and so on – will be centralized through a single interface. For example, you can already ensure that a meeting you are planning in Indico is exported automatically to your personal calendar. Similarly, you can book a subsidiary VRVS session directly from Indico, rather than having to do it in VRVS. We're doing similar things to the e-dial system for audio conferencing, and have a major development under way to integrate the conference room booking system, which will probably become a module of Indico in the future.

In your view, what is the future of videoconferencing at CERN?

We recently set up an H323 service, which is the commercial offering for IP-based videoconferencing. It is an open standard so many manufacturers are making compatible hardware and software. It is also a high-quality standard, which means that the end devices tend to be more expensive than simple webcams. CERN is collaborating with the French National Institute of Nuclear Physics and Particle Physics (IN2P3) in Lyon on this project, and the centralization unit is being housed at IN2P3.

VRVS is today the more dynamic, self-learning infrastructure, and it can support more capacity and many more end devices. This will remain popular in the particle physics community because the cost of joining is lower. But as H323 becomes more commercially widespread the difference could become less significant and more HEP users might be attracted. We are now equipping the conference rooms with high-quality H323-compatible equipment. We are not dropping support for VRVS, just standardizing on a set of equipment that can support both technologies.

The CERN Document Server software has recently been renamed Invenio.

Why this change? And how is the role of CDS evolving?

The CDS software is becoming increasingly popular at external sites, and some of these are now contributing to the development of the software. For example, EPFL in Lausanne is a significant contributor. To reflect this broader development and use we rebranded the CDS software as Invenio. This also clearly indicates that it is a sister product to Indico, which originally meant Integrated Digital Conferencing and also means "to announce" in Latin. Invenio means "to find, to discover" in Latin. We will still refer to CDS as the CERN instance of Invenio.

There has been much development of CDS recently. In particular we have gone through a major phase of consolidation, changing the underlying software to Python and thus refreshing the code base. We see a period of expansion ahead. There has been a strong push from the CERN library and the directorate to make CERN's institutional repository the focus of all information at CERN, and to discourage the reliance on subsidiary storage systems and websites for storing physics and

administration documents.

Of course, there is a separate Engineering Data Management Service (EDMS) system for design and engineering documents. We have carried out an assessment with the EDMS team, and we see clear reasons for having product life-cycle management in EDMS that we do not need in CDS, so merging these two technologies is not on the cards.

With this proviso, the aim now is to ensure that everyone deposits information in CDS, and that all of the data is interlinked and can be found efficiently. As part of the expansion phase ahead we are modifying the architecture for the servers to make them more scalable, and at the same time increasing the server base. The number of documents in CDS is growing at over 100 000 per year, and we have recently absorbed many Web archives into the server.


Finally, a key role for CDS in the future is to give the organization the necessary tools to implement the open-access policy to scientific literature, which the CERN library is spearheading with the director general's support.

Computing featured in this month's CERN Courier

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

Computing News

- EGEE gets down to business
- OSG receives \$30 m for scientific computing collaboration
- KnowARC project gets going
- GRIDtoday honours CERN and e-Science
- W3C to support more languages in voice applications
- Jefferson Lab attracts record numbers to Geant4 workshop
- CNL celebrates 40 years of news



- LHC@home moves its base to the UK

Calendar of events

Feature article

- Symposium focuses on scientific supercomputing

Computing at CERN from the 1990s to the Grid era

We are now reaching the end of our journey through the archive to celebrate *CNL*'s 40th anniversary. Last issue we looked at the Web revolution that took place in the nineties, and this issue we will cover the period from the early 1990s until 2002, when references to the Grid start to appear. Again we will reproduce extracts of articles that were published during that time.

The first Web version of *CNL* was published in October–December 1993. You can access this issue at <http://cern.ch/cnlart/214>. URLs are also provided so that you can retrieve the full stories of the following excerpts.

Evolving CERN computing strategy and changes in CN's structure

“Recent years have seen a revolution in industry’s approach to building and marketing computer systems and in the ways people use them. This has been characterized by a trend away from large monolithic systems whose high cost must be amortized over many users and towards a more distributed, scalable and personalized approach, based on compact CMOS machines manufactured in large volume with the benefits of a vastly improved price/performance ratio...”

David Williams, head of CN division (which preceded the IT department). (October–December 1993, <http://cnlart.web.cern.ch/cnlart/214>)

Interview with Knuth on software tools and techniques

“Donald E Knuth is considered by many to be the world’s preeminent computer scientist... Knuth has also developed breakthrough applications in computer typesetting (TeX and METAFONT) and software development (WEB), and has over 100 publications to his credit...”



Renovation of the Computer Centre in 2003. Half of the equipment in the ground floor room has been removed to prepare for the PC farms.



Above: View of the computer room, building 513, in 1998. Right: The PC farm for the Grid in the Computer Centre as it is today.



Interviewer: What about object-oriented programming? Is it just a current buzzword, or does this approach appeal to you?

Knuth: I’ve always thought of programming in that way, but I haven’t used languages that help enforce the discipline; I’ve always enforced the discipline myself in other languages... The problem that I have with them today is that... C++ is too complicated. At the moment, it’s impossible for me to write

portable code that I believe would work on lots of different systems, unless I avoid all exotic features... But each user of C++ has a favourite subset, and that’s fine. CWEB fully supports C++ as well as C.” (April–June 1994, <http://cern.ch/cnlart/216>)

Crack, a password checking service at CERN

“Passwords are the primary security loophole on a system. If

users have guessable passwords then ‘bad guys’ can use their accounts without being detected. These intruders can cause damage to the individual user and also to other users. The use of security tools like Kerberos or shadow passwords are not sufficient since the ‘crackers’ can still try to guess passwords.

The only really safe solution is to make users aware of the dangers and understand the following:

- why it’s dangerous to have guessable passwords (for them and for CERN);
- how to choose a good password...

However, some guessable passwords will always be used and conscientious system administrators should try to reduce the number of users with a guessable password. This is the purpose of the ‘Crack service’.” (July–December 1995, <http://cern.ch/cnlart/221>)

The “Year 2000” problem: the millennium bug

“The millennium bug is not some exotic insect, but an error that is lurking in computer programs, waiting for the chimes of midnight on 31 December 1999, before emerging to bite us – hard. Even as the champagne corks are popping to herald in the last year of the second millennium, planes will start to fall from the sky, lifts will stop working, banks will stop paying correct interest, and society in general will grind to a halt – or so the story goes.

To illustrate what the fuss is about, consider an early sign of the bug’s presence. In 1993, a lorry carrying a consignment of corned beef attempted to deliver its load to a British supermarket. Corned beef has a shelf life of seven years, but when the supermarket’s computer tried to check this, instead of subtracting 1993 from 2000 to get seven, it

performed truncated arithmetic, subtracting 93 from 00, to give -93! The load was refused (but accepted manually)...” (June–September 1997, http://cern.ch/cnlart/228/art_milb.html)

Remote/mobile computing at CERN

“With the growing number of portable computers, palmtops and mobile phones, as well as the need to keep in touch with CERN not only from home but also whilst on business travel, a coordination activity has been started in IT Division for remote/mobile computing...” (July–September 1998, http://cern.ch/cnlart/232/art_mobcomp.html)

A new mail address for SPAM complaints

“Following a *de facto* standard

in the Internet world a new contact address at CERN has been defined for matters related to unsolicited bulk mail (SPAM). Please direct your questions or complaints concerning SPAM to abuse@cern.ch.

This will allow us to share the work load more efficiently...” (January–March 1999, http://cern.ch/cnlart/234/art_spam.html)

SHIFT software evolution: CASTOR

“In January 1999 CERN began to develop CASTOR (CERN Advanced STORage manager). It is an evolution of SHIFT, which has been used at CERN for all experiments in the 1990s. CASTOR is more scalable than SHIFT and should be able to handle data for the overlapping run of the NA48 and COMPASS experiments, the latter

starting in 2000...” (October–December 1999, <http://cern.ch/cnlart/237/shift-castor.html>)

European DataGrid project demonstrated successfully

“On Friday 1 March, the EU-funded DataGrid project successfully passed the first year review performed by external experts appointed by the European Union.

In late 1990, a CERN computer scientist invented the World Wide Web to facilitate exchange of information between scientists working on different computers, perhaps at different sites. At the EU review, CERN, along with DataGrid project partners INFN/CNAF (Bologna, Italy), CNRS/IN2P3 (Lyon, France), PPARC/RAL (UK), and FOM/NIKHEF (Amsterdam/NL),

demonstrated the first DataGrid test bed. The WWW is mainly aimed at the exchange of information while the Grid is concerned with the exchange of computer power, data storage, and accessing large databases, without forcing users to search for these resources. Once ‘connected’ to the Grid, the end user will see it essentially as one large computer system.

For these reasons, many believe the Grid to be the most practical solution to the so-called data intensive science problem that must be overcome if the computing needs of scientific communities such as processing of physics data from LHC experiments are to be satisfied...” (January–March 2002, <http://cern.ch/cnlart/2002/001/datagrid>)

CERN mail service requires secure protocols

Insecure access to central e-mail servers on the MMM (My Mail and More) service was closed on 2 October. This came after several campaigns to enforce the use of secure e-mail protocols (see the April–May issue of *CNL*).

There are three main changes to the mail service:

- Encryption is required when reading messages from a mail client using the POP or IMAP protocol. Detailed instructions on how to enable this security option for Pine, Mac.Mail and Mozilla/Thunderbird clients are available at <http://cern.ch/mmm/services/Help/?kbid=191040>.

ch/mmm/services/Help/?kbid=191040. CERN clients preconfigured to Outlook XP/2003 are not affected because they use a different protocol.

- Authentication and a valid sender address are required when posting messages from a mail client using the CERN SMTP service. Valid addresses for each person are defined in the CERN Computing Resources Administration (CRA) system. See <http://cern.ch/mmm/services/Help/?kbid=191060> for details of how to check your valid addresses. This page also

explains how to enable another person to send messages using your e-mail address.

- Anonymous posting is still supported for cron jobs, scripts and internal mail servers (with restrictions). These insecure connections are only possible using a dedicated server, cernmx.cern.ch (see <http://cern.ch/mmm/services/Help/?kbid=191010>). Posting is temporarily denied for a machine when it reaches certain thresholds: 3000 mails per day or 50 external recipients (not @cern.ch) per day. Machines

that need to be above these thresholds can be white listed.

People who use Pine should also read “Known issues for Pine mail client users on Linux” on p15.

The CERN Mail team would like to thank all of the users and machine owners who have updated their mail settings to comply with this new policy. These changes will greatly reduce the risk of malicious programs (such as spambots and viruses) sending unsolicited e-mails over the Internet via the CERN mail servers.

The CERN Mail team

Printing of banner pages to cease

The printing of banner pages has been disabled under the new infrastructure to reduce paper consumption. Statistics show that the average print job size is small, and paper savings of up to 20% could be made by not printing the banner page.

As each printer is moved to the new infrastructure, banner page printing will be disabled. The helpdesk can re-enable banner page printing upon request for corridor printers that are shared by several users. We hope that ultimately the printing of banner pages will be enabled on fewer than 10% of printers

registered on the network.

If banner pages have been disabled centrally you can still print them by using Linux. Simply add them to your print job on the client side by adding the “-o job-sheets” option to your lpr command.

Please bear with us if, during the transition period, banner page printing has been disabled on a printer that requires it. If this happens, just send an e-mail to the helpdesk stating the name of the printer and asking for banner printing to be re-enabled.

The NICE team

Discovery Monday profiles the Grid

Some 150 members of the general public attended a Discovery Monday event about the Grid, held at CERN’s Microcosm on 2 October.

The visitors watched a short presentation about the Grid and its relation to CERN, and about the huge data challenge posed by the Large Hadron Collider experiments. Staff from the IT department then gave several interactive presentations.

One presentation about Grid monitoring tools provided a visually impressive sense of the global collaboration behind the EGEE and LCG projects. Another

presentation, by the UNOSAT team at CERN, showed how the Grid can be used to store and speed the provision of vital satellite images to aid workers in natural disaster and post-conflict zones

Flash animations illustrated how the Grid middleware works in practice; and a Grid game, called Amalthea, introduced some key concepts in an educational way. People who completed a three-minute challenge received GridCafé T-shirts. There were several winners, including a few adults.

The IT Communications team

E-mail addresses at CERN

When staff first arrive at CERN they have to set the addresses that others can use to contact them via e-mail. This article will describe the options available.

Choosing an address and mailbox

As soon as CERN users and visitors are registered in the Human Resources (HR) database, they will receive an e-mail address set to “firstname.lastname@cern.ch”, where the first name and last name are taken from the passport. This is automatic and mandatory, so that users can be contacted easily in case of problems.

This address is referred to at CERN as the GEM (Generic E-Mail). However, the name generated from your passport may differ from the name by which you are generally known, so you may wish to change your GEM. This is described below in “Setting the GEM”.

The GEM does not determine where messages will be delivered; for this users must define the PEM (Physical E-Mail). This corresponds to the physical mailbox where incoming mails are sent. The PEM can be set within or outside CERN.

“Defining your PEM” describes how you can set your mailbox.

Setting the GEM

Most people have on their passport two or more first names, and sometimes a composed family name. If we take, for instance, “Martin John Smith Wilson” (i.e. first name “Martin John” and family name “Smith Wilson”), the GEM address that will be created automatically will be: Martin.John.Smith.Wilson@cern.ch. However, people might know you as “Martin Smith”. If so, you can modify the automatic registration so that people can contact you at the address Martin.Smith@cern.ch. The first name must be redefined to “Martin” and the family name to “Smith”.

First name

Users can redefine their own first name using the Computer Resource Administration (CRA)

interface at <http://cra.cern.ch>. You will need to provide your AIS log-in password the first time you access CRA.

Select the section “CRA/End-Users option” then “My details” to see the data registered for you. You can amend your GEM by adding a nickname that will replace the first name. For example, by giving the nickname “Martin” the GEM will become: Martin.Smith.Wilson@cern.ch.

Family name

Users cannot redefine the family name themselves, instead they must contact the CERN User Registration team (at User.Registration@cern.ch) who can set a synonym. In the above example they can set it to “Smith”, and the GEM will become Martin.Smith@cern.ch.

People who marry while at CERN may also find it useful to change the family name. If they want to be contactable by both their original and married name, they can ask the registration team to also set a synonym for their married name. Similarly, divorcees can request a synonym for their new name.

Defining the PEM

The PEM is the address that corresponds to your physical mailbox, i.e. this is where messages sent to your GEM (ending with @cern.ch) will be delivered. You can choose whether or not to have a mailbox at CERN. This is not mandatory and you may prefer to use a mailbox at your home institute or a private mailbox. If you have a CERN mailbox, the PEM must be internal to CERN and the messages sent to your GEM will arrive there. If you do not have a CERN mailbox then the PEM will be external to CERN, and the messages sent to your GEM will arrive in your external mailbox.

If you want a CERN mailbox you will be given an account on the CERN mail server and your PEM will become username@mail.cern.ch (where “username” is your login id on the CERN mail server). If you do not want a CERN mailbox then you will have

to define an external PEM such as user@fi.infn.it or user@hotmail.com.

If the PEM is external, it must be set by your administrator before he or she creates your other CERN accounts. With an external PEM you will not be allowed to have an account on the CERN mail server. If an account has been created it will be deleted automatically during the night by the CRA system.

If you have a CERN mail account but think you won't need it because you use an external mailbox, then you should set the PEM to your external address. Automatically your mail account on the CERN mail server will be removed by the CRA system.

How to forward messages

All messages to your GEM are delivered to your CERN mailbox, if you have one. However, you can choose to forward incoming e-mails to another address. You can do this via the “Valid e-mail address and forwarding” link on the MMM (My Mail and More) website (<http://cern.ch/mmservices>). You can set one or more external addresses to forward incoming mails, or can have messages delivered to both the forwarding address and the CERN mailbox.

Spam

Users should be aware that if they have defined a PEM address outside CERN, or if they forward messages to an external address, then spam will not be filtered and all mails including spams will be sent to the PEM mailbox or forwarded. It is up to the mail service on the external site to provide a spam filter.

Phone book

In xwho, the CERN phone book on the Web (at <http://consult.cern.ch/xwho/people>), the GEM is displayed in the preferred form requested by the user.

To refer to our example, xwho will display the address Martin.Smith@cern.ch. However, xwho also provides a link (“More info about valid mail addresses”)

where it is possible to see the PEM and other GEMs that may exist. The PEM is the first entry in the list. In our example, supposing that Martin Smith has an account on the CERN mail server (login id “msmith”), you will see the following list of valid e-mail addresses:

- msmith@mail.cern.ch. This is the PEM to which all GEMs are pointing. The address should be changed whenever you change the system (server) for reading your mail.

- Martin.Smith@cern.ch. This is the GEM that is guaranteed to always work – it will always point to the person's PEM, even if the physical address is changed. This is the address that should be put on a business card and otherwise published. Once established it is strongly recommended not to change this address again.

- Martin.Smith.Wilson@cern.ch. This is also a GEM that works at the moment but may stop doing so if a new person with the same name is registered at CERN.

- Martin.John.Smith.Wilson@cern.ch. This is another GEM that works at the moment but may stop doing so in the future.

If Martin Smith did not have an account on the CERN mail server, but instead had defined an external PEM set to, for instance, msmith@rl.ac.uk, then the first line in the list would be: msmith@rl.ac.uk – the PEM to which all GEMs are pointing. The remaining entries in the list would be as above.

Accented characters

At the moment and for some time to come the HR database cannot store “accented” characters (e.g. é, ü, ä, ö, â, ñ), so GEM addresses cannot include such characters. The Web phone book can display these characters, but only at the second level, which you reach via a click from the xwho search output. The procedure is described at <http://consult.cern.ch/xwho/help/correctnames>.

Nicole Crémel and Roger Woolnough, IT/UDS

Virtual servers cut hosting costs and offer flexibility

Virtualization is back. However, today it is more the renaissance of a dying technology than a brand new concept. Like the videophone or the new Geneva tramway, this technology has existed for decades.

The virtual computer is an old dream, a long story associated with the evolution of the computer. The goal has always been to make the software independent from the hardware, enabling the hardware to evolve without having to rewrite the software. We all know what it is to suffer from service migrations, server reinstallations and hardware renewals, which generate inconvenient interruptions to our work.

However, virtualization has always faced performance problems: the need for better performance in terms of speed and development time has always been opposed by the need for portability and investment preservation.

Historical perspective

Between 1977 and 1979 Unix tried to address the issue of virtualization with the C language: "Write it in C, it will run on any Unix." But just a few years later it was shown that the C language did not solve all of the problems.

In the 1980s the IBM 390 system had specialized circuits in the CPU that enabled it to virtualize itself. The virtual machine idea was there: IBM's Virtual Machine Operating System provided each user with his or her own independent virtual computer.

In 1995 the Web generated another virtual machine: the Java Virtual Machine. The idea was to have a virtual machine and an environment that was independent of the underlying operating system. This time the slogan was: "Write once, run everywhere." Again all problems were not solved by Java, but this

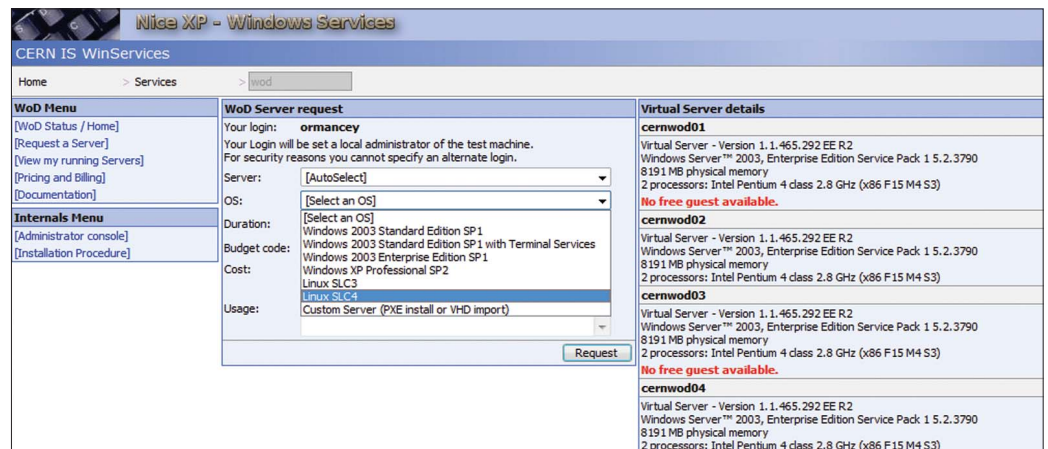


Fig. 1: Users of the custom Windows server hosting service can request a brand new server in several clicks.

time it may have been due to commercial interests rather than to technical difficulties.

So what's new?

Today virtual technology has become standardized: the Intel PC is now a consolidated standard. The Intel Virtual PC has become the new virtual machine base layer. It can run any operating system supported by the Intel PC, from Windows to Linux, and possibly Mac OS. It can run on any PC processor vendor (Intel or AMD) and on any processor model (Itanium, Pentium 4, Opteron and so on).

Today's emulation overhead is small and the emulation of PC on PC is very efficient. Also, Intel and AMD processors now integrate a dedicated virtualization technology layer to enhance the performance of emulated machines.

A case for virtual servers

CERN's NICE custom servers are a good candidate for virtualization. At the moment the IT/IS group provides a custom Windows server hosting service (in the NICE environment) for groups and departments that require an "out of the box" Windows server solution (see <http://cern.ch/Win/Help/?kbid=251010>).

These custom Windows servers are provided with a high-availability service level agreement: servers are hosted in the CERN Computer Centre, connected to uninterruptible power supplies and monitored around the clock. Daily backups, hardware and operating system maintenance, security scans and patches are all transparent to the customer.

With this kind of solution the customer is free to focus on the applications, rather than on the server that is hosting them. Most of the time customers are unwilling to share the server with others and are ready to pay a kind of rental.

This service is popular. IT/IS receives several requests each month from sources such as LHC controls and experiments (like ALICE), technical services and video streaming services.

Today more than 60 servers run the custom Windows server hosting service, and some weaknesses have been found that could be solved using virtualization. Installing and maintaining physical servers is time consuming, and it requires a management overhead for logistics and resource planning. The space in the CERN Computer Centre is also a scarce resource that cannot be extended

infinitely. IT/IS group has also noticed that some of the servers are underutilized, with only about 2–3% of the CPU in use.

Managing hardware

Virtualization creates a clear separation between the management of hardware and server (software). This management could even be performed by different teams.

The tasks required to manage hardware include basic maintenance, but mostly involve supporting a large pool of servers. The management team has to ensure that enough server hardware is available to satisfy the global demand for CPU and storage.

The team allocates server images to machines in the pool, manages server configuration, and can consider possible optimizations and automations (for example, an automatic reallocation to different hardware according to past performance requirements).

Advantages of virtual servers

Managing virtual servers is much simpler: "installing a server" becomes "loading a virtual machine image". And unprecedented automation can be achieved.

The IT/IS group has started

Hostname	Guestname	OS	Start / End Date	Status	Action	CPU
W004-001	VM04-001	W2003SESP1 [Export] [Import]	7/4/2006 12:39 PM 7/11/2006 12:39 PM [Extend]	Running	[StandBy] [Resume] [Reset] [TurnOff] [TurnOn] [Admin Console]	

Fig. 2: The server's Reset and TurnOn/Off buttons are links on a website.

Fig. 3: A website gives access to adapters, a CD/DVD reader and so on.

Table 1: Read/write speed on emulated drives compared to physical host RAID1

	Physical	IDE emulated	Dynamic	Static
Write (Mbyte/s)	56.11	30.27	55.43	
Read (Mbyte/s)	54.16	40.07	42.16	

automating the Microsoft Virtual Server 2005 software by providing a Web interface for users of its custom server service (see <http://cern.ch/WinServices/Services/WoD>).

The user can request a brand new server in just several clicks: Select an OS, Duration, Budget code and Usage, then click Request (figure 1). Ten minutes later the virtual server is available and ready for use. The user will automatically be added to the local administrator's group for Windows servers or to the sudo users list for Linux servers.

A budget code is required to make sure that users don't abuse the virtual server facility. Although a new server is just a few clicks away, computer resources at CERN are not free.

Once the server is installed and ready the user can interact with it without leaving the office: the Reset and TurnOn/Off buttons are links on a website (figure 2). The user can export and import the disk image through this interface: it can be stored for use on the user's server after it has been configured.

A Web interface provides access to the server console, enabling BIOS settings to be changed and various configurations to be made. The server's physical configuration can also be edited using a website that provides access to a network card, CD/DVD reader, floppy disks and so on (figure 3).

Virtual servers on demand

The virtualization at CERN has been automated to fit the organization's generic needs: various operating systems can be requested, like Windows 2003 classic, Windows 2003 with IIS (Web server), Linux SLC3 and SLC4.

An empty image can also be requested: a virtual server with no operating system will be provided and the customer can install manually his or her own operating system. This can be done in two ways:

- The user can mount a CD/DVD image and boot from it to run the installer, or press F12 to boot from the network (PXE) and start various network installations.
- Taking advantage of disk

images used in virtualization, the user can import a disk image to this virtual server guest. This disk image can be generated from a physical computer (any software that is compliant with VHD file specification can achieve this, like WinImage shareware), or come from a previous virtual server on which the user exported the disk image using the provided Web interface.

When a customer requests a new server the resource is taken from the pool of available hardware: multiple and different operating systems can be hosted in the same box. The server is available in 10 minutes (the time it takes to copy the image then join the CERN domain for Windows OS, or install updates for Linux OS).

Sharing hardware leads of course to an important cost reduction: the IT/IS platform now hosts three virtual servers (guests) for each physical machine (host).

Performance

Users of the virtual servers have not noticed any difference in service. Of course an application that is input/output and CPU intensive could suffer from virtualization, but classic applications do not.

Today's service emulates servers with one CPU in 32 bit mode, running at the same speed as the host (2.8 GHz or more). A CPU can be preallocated – one guest can reserve 100% of one CPU. By default the three guests share the host's two CPUs according to their needs. The memory allocated is 2 GB per guest.

Concerning input/output, the emulated disk drives can be IDE or SCSI, and can be dynamic or static. A dynamic drive will use only what it needs on the physical disk, compared to the static disk, which will use the declared emulated size. Table 1 shows the read/write speed of dynamic and static drives compared to their physical host.

The idea, however, is to use centralized file space, like DFS or AFS, to avoid storing important data on virtual server guests. This helps to replace a guest server more quickly, and enables data to be accessed easily from multiple guests.

Concerning the network,

emulated cards have their own generated MAC address and share the physical host Gigabyte network card.

Servers on demand

The IT/IS group anticipates that it will provide more server types, with various combinations of operating systems and applications. Requests for custom server types are also expected, where users will import (and export) their own server images. A physical server can be virtualized by simply building an image of its hard drive(s).

A "server on demand" service must be able to satisfy the following requests:

- "I need 20 servers with this image for one month."
- "I need an image for this server replicated 10 times."
- "I need more CPU/memory for my server."
- "I need a test environment, OS version $n+1$, to which I can migrate my current production services."

Batch systems

We can also imagine the future of batch systems. Instead of sending a piece of code to a pool of batch servers, the batch creator could virtualize his or her own machine (desktop, development environment, etc.), and send this virtualized image to a pool of virtual servers. The user would then run the batch in his or her preferred mode.

Running a virtualized machine could also prevent batch scripts or programs from creating problems on the batch server, because all processes are restricted to their own virtual machine.

Conclusion

Server virtualization is a strategic direction for managing servers at CERN: it reduces the cost of hosting and managing servers while providing more flexibility for users. Managing a pool of identical servers simplifies operating procedures, and virtualized guests increase the range of operating systems and applications while remaining independent of hardware.

We can also expect virtual servers to have consequences for traditional batch systems, providing a more robust service.

Emmanuel Ormancey and Alberto Pace, IT/IS

Getting Grids to work together: interoperability is key to sharing

In recent years a number of Grid projects have emerged to help coordinate institutions and enable Grids. Today we face a situation where most of these projects, many of which have a strong regional presence, have slightly different middleware.

Establishing interoperability between Grids is vital to bridge these differences and enable virtual organizations to access resources at the institutions, independent of the Grid project's affiliation. Without Grid interoperability, collaboration would be artificially limited to one Grid or the collaboration would have to create multiple virtual organizations and manage the diversity itself.

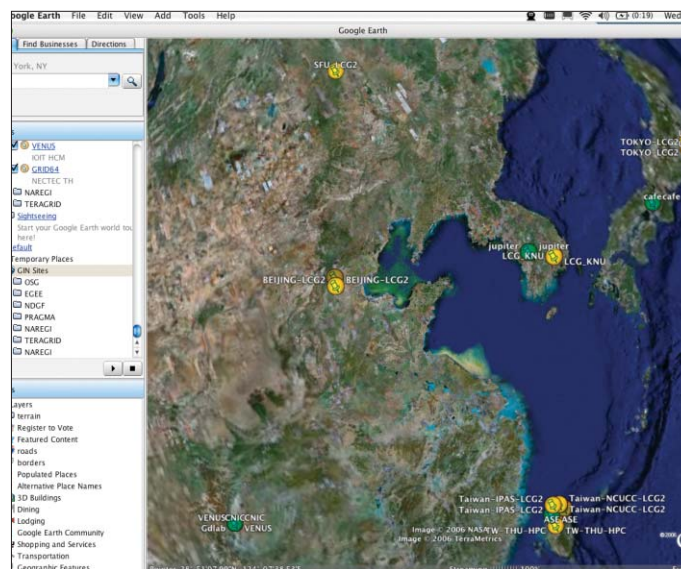
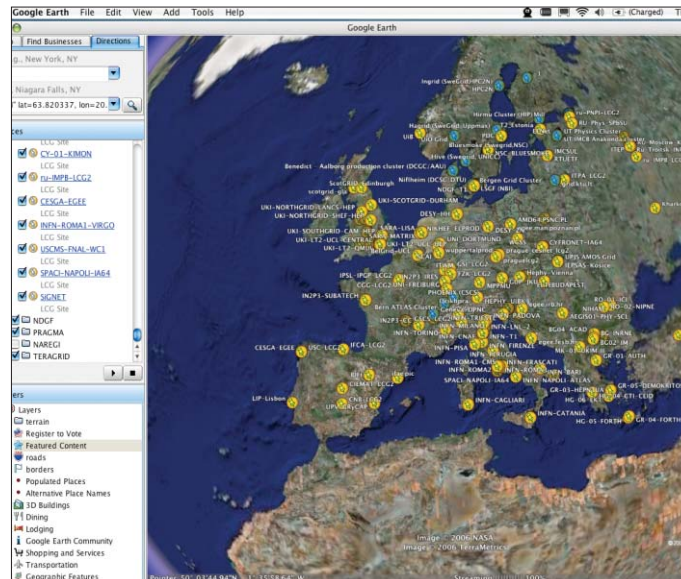
Anatomy of the Grid

To understand the nature of this problem it is worth stepping back and comparing the situation today with the original concept of a Grid. In the influential paper *The Anatomy of the Grid* [1], a Grid is defined as being "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations".

In this statement there are three fundamental entities: resources, institutions and virtual organizations. A virtual organization is a group of users from multiple institutions who collaborate to achieve a specific goal. An institution is an administrative domain and has complete control over the resources within its domain. Institutions support a virtual organization and hence allow users, who may belong to different institutes, access to the resources.

As each institution has control over its domain, each institution may have different systems and policies. To overcome this heterogeneity, Grid middleware is used to provide an interface at the boundary of this administrative domain.

Grid middleware follows the



These maps showing Grids in Europe and Asia use a Google Earth interface that was developed by Imperial College London [2].

"hour glass" model. At one end there is a diverse set of resources and at the other end there are many virtual organizations that have their own applications. The applications can gain access to the heterogeneous resources through a small set of well defined interfaces. As different Grids have their own middleware and policies, they can also be seen as different administrative

domains. In a sense, the challenge of Grid interoperability can be viewed as a problem analogous to that of users accessing resources at different institutions, but now with virtual organizations accessing resources on different Grid infrastructures.

Understanding interfaces

How can interoperability between Grids be achieved? The first

thing to do is compare the interface at the boundary of the administrative domain for both Grids. Once this has been done, it should be possible to create an interoperability matrix between two infrastructures. This matrix will reflect the similarities and differences of the interface. Not only is it necessary to compare the interfaces, it is also important to understand how the interfaces are used. Once the differences have been understood, steps can be taken to overcome these differences.

Common interfaces seem to be the most straightforward approach. However, with the absence of standards, which interface should one choose?

As a Grid infrastructure has invested heavily in one interface, it may be difficult to move to another interface. So although a common interface is the ideal solution, reaching agreement on which interface to use and the deployment of a production quality implementation across all infrastructures will take time.

Adapters and gateways

In the short term, adapters and translators can be used in the higher level services so that the software can work with both interfaces. Adapters bridge incompatible interfaces, and translators convert information to a format that another system can understand. This approach requires some parts of the Grid middleware to be modified, but it does mean that it will be possible to use the existing interfaces. This method will enable interoperability to be achieved without having to modify the existing interfaces, which means the institutes will not be affected. Where and how the adapters and translators are used will highlight areas that need standardization.

If it is not possible to modify the higher level services, then

gateways can be used. A gateway is a bridge between Grid infrastructures. It uses the same technique as adapters and translators but the gateway is a specific service, it is not built into the middleware. The problem with Grid gateways is that they can be a single point of failure and also a scalability bottleneck. As such they are only useful as a temporary solution.

Grid operations

Once technical interoperability has been achieved, it is important to start looking at Grid operations, which cover everything that is needed to operate a Grid infrastructure. This includes service monitoring, user support, resource accounting, problem resolution and so on – issues that are all procedure oriented.

The support teams within the different Grids may rely on different software tools but it is not necessary to harmonize these tools. However, it must be ensured that the tools from one Grid infrastructure will work on the others. The procedures used on each Grid infrastructure need to be analysed to ensure that the necessary operations can still be carried out with the additional institutions and

virtual organizations. For example, ways to route trouble tickets between Grid operations centres need to be investigated.

EGEE activities

Grid interoperation is usually a bilateral activity between two Grid infrastructures. One of the first interoperation activities was between the infrastructure used by the Enabling Grids for E-science (EGEE) project and the Open Science Grid.

The initial analysis showed that the middleware used by both infrastructures was similar, and for this reason both infrastructures decided to use common interfaces. After an initial proof of concept was carried out in January 2005, the changes needed were integrated in the middleware stacks for the respective infrastructures. It took about six months for these changes to be included in the official software releases and rolled out across the infrastructures in August 2005.

After this work was done, discussions moved towards Grid operations. Policies had to be aligned and modifications were needed in the operational tools. After about another six months, by January 2006, the Grid infrastructures were seamlessly

interoperating and virtual organizations were successfully using both infrastructures.

There are now several ongoing bilateral activities. The activity between EGEE and the Nordic DataGrid Facility is well under way and is trying to use a combination of gateways alongside adapters and translators to achieve interoperation. The results of initial testing look promising and discussions on Grid operations have begun. Other infrastructures where EGEE is involved in bilateral activities include NAREGI, Unicore and ChinaGrid.

Moving forward with GIN

Within the Open Grid Forum, the Grid Interoperability Now (GIN) community group has been trying to build upon these bilateral activities. The GIN group is a focal point where all the infrastructures can come together to share ideas and experiences on Grid interoperation. It is hoped that each bilateral activity will bring us a step closer to the overall goal of a uniform Grid landscape. The recent achievements of the GIN group will be demonstrated at Supercomputing 2006 at Tampa in November.

Grid evolution

To achieve interoperation, interoperability between middleware is not the only issue. As Grid infrastructures evolve, different middleware will be used. Of primary importance, then, is to ensure that even if the middleware evolves, interoperation is maintained. Using different Grid middleware stacks is a major obstacle but the problem can be overcome.

Production implementations of real Grid standards would go a long way to help harmonize the Grid middleware stacks and help interoperability. However, the current Grid paradigm is “a Grid of Grids” – different Grid federations working together to provide a seamless Grid infrastructure. As a result, even with technical interoperability assured, a truly federated Grid will bring a whole new set of operational challenges.

References

- [1] Foster *et al.* 2001 *Intl J. of High Performance Computing Applications* **15** 3 200–222
- [2] Maps generated by Google Earth (<http://earth.google.com>) using configuration file at <http://lfield.web.cern.ch/lfield/gin.kml>.

Laurence Field, GD/IT, CERN

CERN openlab student reports on LCG experience

Over the summer of 2006 I was a CERN openlab student with the LHC Computing Grid (LCG) project, cosponsored by my home institute, the University of Milan. My supervisor at CERN was Birger Koblitz of the IT/PSS group, and my project was to develop two applications to help users to query metadata catalogues on the Grid: a generic browser for metadata and a browser for the LHCb bookkeeping database.

To find the files of interest in a large data Grid, an efficient mechanism to discover and query information about the stored data is needed. This is achieved by attaching descriptive attributes (metadata) to files and by listing this information in catalogues. The Enabling Grids for E-science (EGEE) project's gLite middleware provides this service through the AMGA (ARDA



View from the top: openlab student Danilo Piparo on building 40 at CERN.

Metadata Grid application) metadata catalogue.

The AMGA generic browser is a generic metadata graphical user interface that enables an interactive exploration of the

metadata schema and entries. The basic idea was to help the user to browse the metadata catalogues on the Grid with a familiar explorer-like interface.

The LHCb bookkeeping

database browser is a specialized version of the generic one. It was developed for the LHCb experiment to ease the access to the logging and bookkeeping database. This software was designed so that it can be a part of the Ganga GridUI of Atlas and LHCb, and I integrated it into version 4.2.0-beta5 of Ganga.

My work experience at CERN culminated at the end of September with my participation in the EGEE'06 conference in Geneva, where I presented the results with a demo of my software and a poster.

I very much enjoyed this summer at CERN. I worked with a group of professional developers in the context of LCG and learned a lot from working in this environment. The openlab student programme was an enriching experience for me.

Danilo Piparo, University of Milan

EGEE'06 debates future of Grid

The EGEE'06 conference, held during the last week of September in Geneva, Switzerland, brought together some 700 representatives from many Grid projects. These included the Enabling Grids for E-science (EGEE) project itself, as well as users and major players in the business sector. The large turnout made it the biggest Grid conference in Europe since the Fifth Global Grid Forum in Edinburgh in 2002.

"The community spirit seen at EGEE'06 was overwhelming, both within the project consortium and with the many other organizations and projects we work with," said Bob Jones, the EGEE project director. "The achievements demonstrated this week show the quantum leap that Grid computing has taken in the last few years, from an interesting idea to a reliable technology that has been adopted by a broad range of scientific and research communities."

Many Grid projects took the opportunity to fine-tune their developments and formulate plans for the future of Grid technologies. A record of more than 30 Grid projects came together, in areas ranging from regional infrastructures to domain-specific applications, highlighting the strong international collaboration in this field.

Business role for Grid

Industry's adoption of Grid technology was another highlight at the conference. There was a presentation dedicated to business aspects of the Grid, and it was announced that three companies have signed up as the first EGEE business associates. NICE Italy, Platform Computing and PricewaterhouseCoopers will work with EGEE to make the distributed computing infrastructure of the Grid more user friendly, effective and secure in an industrial context.


Technology transfer is a key goal of projects funded by the European Commission, and, as the flagship Grid project in the Sixth Framework Programme,



The EGEE'06 conference in Geneva attracted some 700 participants.



The UK-CERN collaboration GridPP demonstrated its 3D Grid monitor.



The demo and poster session encouraged some lively discussions.

EGEE will pave the way to making Grids the technology of the future for business in all sectors.

Grid applications took centre stage on Tuesday, with a demo and poster session showing examples of the ever growing number of application domains and individual research groups using the processing and data-storage capabilities offered by the Grid. Users in more than 150 virtual organizations from fields as diverse as biomedicine, Earth sciences and high-energy physics are now using the EGEE infrastructure. The two best demonstrations of

applications shown at the conference received a prize.

"The choice for the best demo was difficult since there were several very nice ones, so we decided to share the prize," explained Alan Blatecky, Deputy Director of the Renaissance Computing Institute and head of the EGEE'06 selection committee. "The climate data analysis shown on EGEE by Kerstin Ronneberger and her colleagues from DKRZ [the German Climate Computing Centre] clearly demonstrated how the Grid can speed up analysis of data and modelling. On the other hand,

the Dashboard for the ATLAS and CMS LHC experiments by Julia Andreeva and her team at CERN demonstrated a great device for monitoring the status of Grid resources and the applications using them."

The interoperability of Grid middleware was also discussed. The EGEE middleware distribution gLite, which was first released for production use in May, is now deployed by about 80% of the sites connected to the EGEE infrastructure, making it the main middleware distribution in use.

EGEE is working closely with projects such as Condor, Globus, Naregi, the Open Science Grid and UNICORE, as well as Grid standardization bodies such as the Open Grid Forum (OGF) to ensure interoperability of the different middleware implementations now available and to work towards universally adopted standards for Grid middleware.

Establishing standards

The challenge faced by EGEE and collaborating projects is to pave the way for a long-term sustainable Grid infrastructure, ensuring that the resources and knowledge developed during its lifetime are available to researchers in Europe and beyond in years to come.

"A necessary requirement for the pervasive adoption of Grid technology will be to break through non-standard barriers," said Mark Linesch from Hewlett-Packard and the president of the OGF. "We need to bring diverse communities together to reduce overlap and align technical directions and requirements, so that we can arrive at mature Grid standards and best practices."

The next major EGEE event, the EGEE User Forum, will be held jointly with the OGF20 meeting in May 2007 in Manchester, UK. This combined event will further strengthen the links between EGEE and the OGF and will bring users and standards bodies together to ensure that the future of the Grid is complemented by the establishment of key standards.

Hannelore Hämmerle, IT/EGE

Digital library solutions employ Grid technology

During the recent EGEE'06 conference, CERN openlab and the EGEE project continued their relationship with [rezonance.ch](http://www.rezonance.ch), organizers of the First series of business networking events.

The First event on digital libraries and e-archiving attracted some 130 people, including experts and members of the international and business communities from around Geneva.

Tim Smith, group leader of IT/UDS, discussed the

challenges that CERN and the scientific community face in archiving their data, and the strategies adopted to solve these issues, which include digital libraries.

Yvonne Grandois, the library coordinator for the World Health Organization, then talked about the WHO's Global Health Library project. This initiative aims to provide global access to medical and health information, both on traditional media such as paper

and in electronic formats.

After these presentations about complex archiving needs, some novel solutions to these problems were offered. Steve Hagner from Oracle Corporation discussed his firm's digital library solutions using database and Grid technology. Pasquale Pagano, of the DILIGENT and DELOS digital library projects, spoke about the evolving European vision for digital libraries and the new work

carried out by the DILIGENT project in linking digital libraries to Grid technology.

Despite the late hour the talks generated a lively panel discussion, ranging from the emergence of digital paper to the open-access movement.

All of the talks are available online at <http://indico.cern.ch/sessionDisplay.py?sessionId=146&slotId=0&confId=1504#2006-09-28>.

Owen Appleton, IT/EGE

Openlab event promotes entrepreneurship

A Grid Entrepreneurship Day was held in the IT department on 26 July as part of the CERN openlab student programme. The aim was to raise awareness among the summer students and young researchers involved in Grid projects of the challenges and opportunities that high-tech entrepreneurship can offer.

Bob Jones, the director of Enabling Grids for E-science (EGEE), kicked off the event, emphasizing that part of EGEE's mission is to promote the uptake of Grid technology by industry.

Nathan Hill and Alex Efimov of the Particle Physics and Astronomy Research Council's (PPARC's) Kite Club gave a series of general talks about entrepreneurship, with advice, examples and a healthy dose of reality and humour. The audience will remember that FFF stands for the typical first-round investors in a new venture: family, friends and fools. The PPARC representatives also

presented the first results of a PPARC-sponsored analysis of the industrial potential of EGEE.

Two case studies were presented of European start-ups involved in Grid and cluster technology. The first case study was a French initiative, Kerlabs, that has an interesting business model for cluster computing based on open-source software. Pascal Gallard, the chief technical officer of this early-stage start-up (it should be established officially by September), explained how the host Institut National de Recherche en Informatique et Automatique played a constructive role in stimulating this spin-off.

The second case study was a Danish firm, MESSTechnologies, which has developed several Grid-related products. The chief executive officer, Søren Nielsen, emphasized that good ideas are not enough – you need to find customers, which can be

challenging when dealing with new technologies. He demonstrated their latest offering, a broadband computer – essentially a thin client for running open-source PC software over broadband connections – which caught the audience's attention.

Presentations were also made by representatives of regional organizations tasked with promoting spin-offs. The French and Swiss representatives, Bruna Carchia and Daniel Loeffler, made a joint presentation to emphasize the close cross-border collaboration that is going on in CERN's immediate neighbourhood. A study of some 20 spin-offs from CERN, commissioned by the representatives, found that most were set up by CERN users, not CERN staff. A significant majority of Italian entrepreneurs was another surprising result.

Antti Heikkela emphasized the approach that Finland has

taken to ensure an effective transfer of technology from CERN and other large facilities. In addition, Beatrice Bressan of CERN's TT group presented the CERN approach to knowledge and technology transfer.

The take-home message for the nearly 20 openlab students and 30 other people from CERN and the region, was that starting up a technology company is not something to do if you want to get rich, but it can be extremely stimulating if you want to be your own boss. The students were also reminded that starting young is a big advantage.

PPARC was a special sponsor of this event. The programme, video and PowerPoint presentations can be found on the openlab website at http://openlab-muinternal.web.cern.ch/openlab-mu-internal/News-Events/Event_pages/06-02_Grid_Enterpreneurship_Day.asp.
Neasan O'Neill, GridPP; Francois Grey, IT/DI

CERN is host to EuroPython 2006

Software developers from around the world gathered in Geneva, Switzerland, on 3–5 July for EuroPython 2006, the fifth annual conference for the Python programming language. The event was hosted this year at CERN, the birthplace of the Web.

"Python keeps growing each year, and EuroPython keeps getting better as well," said



Benedikt Hegner, host and a lead organizer of the conference. "This year saw Guido van Rossum,

Python's creator, move to Google with other key Python developers. Hosting EuroPython at CERN is a prestigious step for EuroPython as well."

The keynote speakers at the conference were Guido van Rossum and Dr Alan Kay, an originator of object-oriented programming and pioneer of graphical computing.

The talks, given in eight tracks, covered the use of Python in science and industry, business and applications, agile development, games and

entertainment, Web frameworks, teaching and more.

The organizers followed the success of earlier years by providing lightning talks, development sprints and "birds-of-a-feather" gatherings – three favourites for fast-paced discussions and networking.

The conference slides and papers can be accessed from the Conference Timetable at www.europython.org/sections/tracks_and_talks.

For news of EuroPython 2007 see www.europython.org.

Printing at CERN: procedures for each operating system

The main aims of the new printing service at CERN are to reduce the amount of home-written software used in the generic computing infrastructure, and to realign CERN with established printing standards. The benefits of these changes will be to reduce costs and to make printing simpler for users, especially visitors, who will no longer need access to special client software.

The new service is based on client-side printing, where each print job is preformatted on the client machine and the print server provides a spooling mechanism to the printer. This requires an appropriate, correctly configured driver to be installed on the client computer.

For Windows clients, the print server will distribute the appropriate driver and default settings. For Scientific Linux

CERN (SLC) 3 and 4, the Printing Wizard software has been modified to install so-called Foomatic drivers on the client (www.linuxprinting.org/foomatic.html).

This article explains how to print from different operating systems on printers managed by the central printing service. To verify the status of a particular printer, visit <http://cern.ch/WinServices/Services/Printers>.

Windows computers (NICE)

Under the new printing system, searching for and installing a printer on a NICE computer is integrated in the operating system and there is no need for a CERN Printer Wizard Windows application. All supported Windows versions have built-in search-for-printer capabilities. The following instructions apply to the NICE XP computer but the

procedure for the NICE 2000 machine is similar.

To add a new printer, go to the Start menu, select Control Panel, and click Printers and Faxes. You will see a list of printers that have already been installed (if any). Click the Add a Printer link on the left side of the window. The printer wizard will start. Select the option: A Network Printer or a Printer Attached to Another Computer, which appears in the first window. Choose Find a Printer in the directory and click Next. You should see a window where you can specify your search criteria. Figure 1 shows how to search for all printers located in building number 31 with a model name that includes the "2500C" substring, and that can print in colour.

Once you have found the printer, just double click on it.

The appropriate driver will then be downloaded from the server and installed on your computer. The default printing preferences will be set as well. You can then verify that your printer is on the list of installed printers (figure 2).

By right-clicking on the selected printer you can set it as the default printer and also choose any printing preferences. Depending on the driver, you can select several options. The most popular are settings for double-sided printing or for default paper size. Once you apply changes, these will become the default printing preferences used by all other applications. You can still overwrite these settings if necessary before each print job.

There is an important difference between a printer configured with the new system and with the old CERN Printer

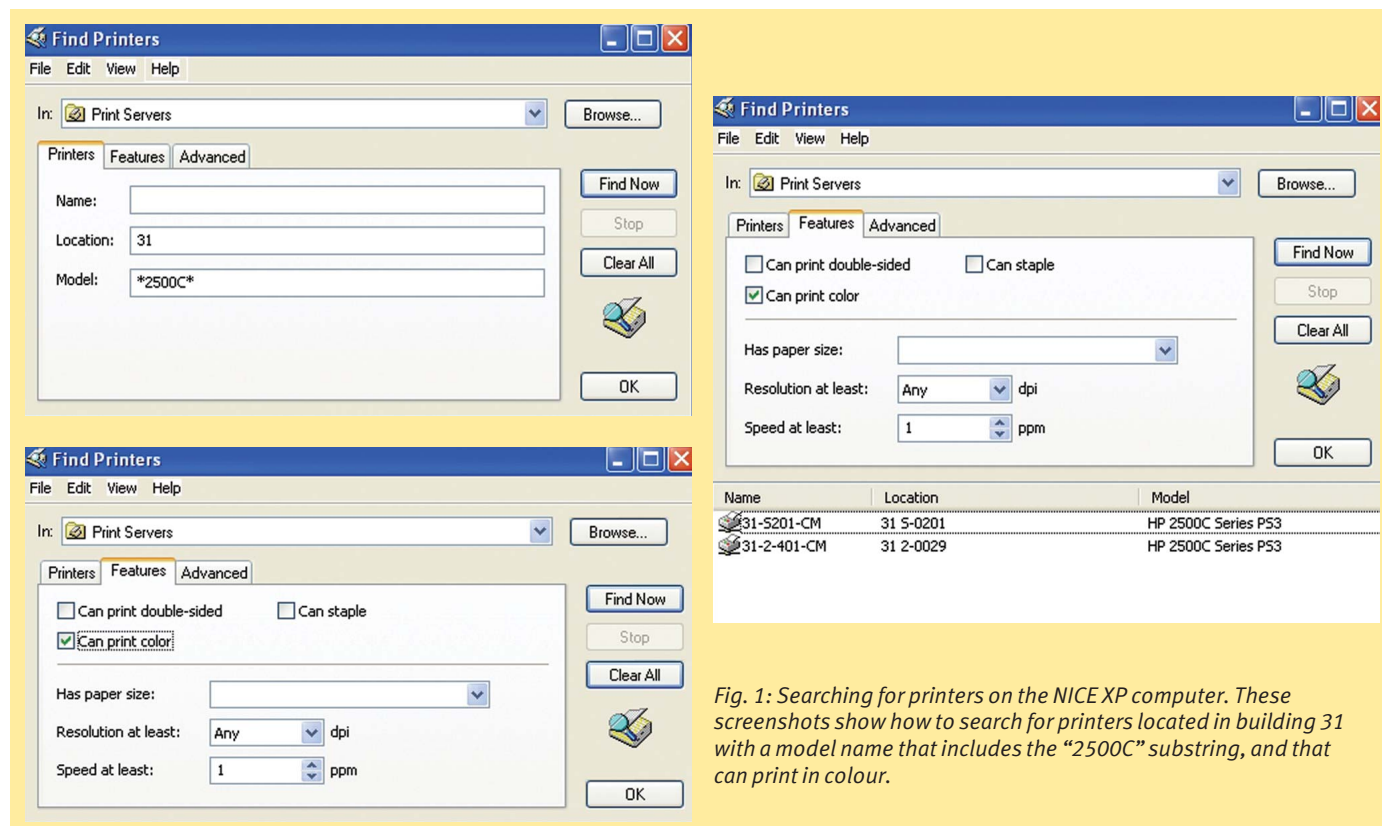


Fig. 1: Searching for printers on the NICE XP computer. These screenshots show how to search for printers located in building 31 with a model name that includes the "2500C" substring, and that can print in colour.

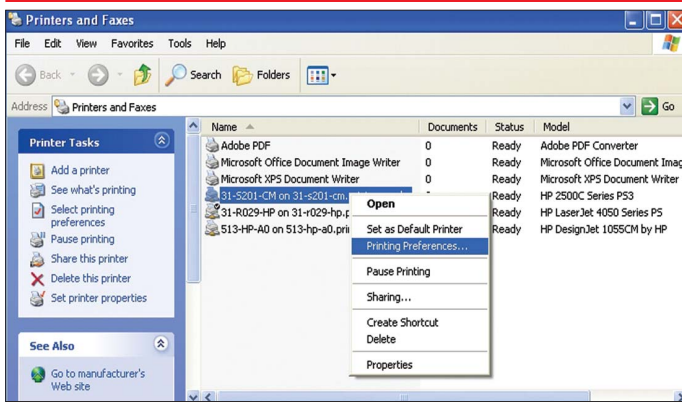


Fig. 2: The list of printers that are configured for the NICE XP computer.

Wizard. Under the new system, all printers on your computer are installed in the user's profile. Now when you install a printer, it is installed only for you and not for other users of your computer. This means that a user does not have to be a local administrator to install a printer.

Scientific Linux CERN 3 and 4

Before users of SLC3 and SLC4 can print anything they have to configure the printers on their computer. They can do this in two ways:

- By using the CERN Printer Wizard graphical interface (figure 3). This tool enables you to search for printers by building and by some general capabilities. You can then install one printer at a time and configure default preferences such as double-sided printing (figure 4).

- By running a command line configuration tool. For example, type "cern-config-printers -b 31 -i" to install all printers located in building 31. Please type "cern-config-printers" with the "-h" parameter to get additional information.

In most cases you will probably print documents from the application itself. For instance, you can print a PDF document by selecting File / Print in Adobe Acrobat Reader. In that case you should be able to select some printing properties from the graphical interface, assuming that the correct driver for a given printer is installed on your distribution of SLC.

To print from the command line you should use the "lpr" command. If you print standard A4 documents on the paper in the default tray, you only need to provide the printer and file names, including the document

you want to print: "lpr -P31-r029-hp mydocument.ps".

If, however, you want to disable double-sided printing or print on transparencies that are loaded in a tray that is not the printer's default, you need to provide additional parameters. Detailed documentation is available about each SLC3 or SLC4 machine at http://cern.ch/linux/documentation/cups-1.1/index.html#STANDARD_OPTIONS. You can also find information on the Linux support pages at <http://cern.ch/linux/scientific3/docs/printing.shtml>.

As an example, if a user wants to print a document on transparencies that are loaded in the upper tray of an HP Business InkJet 2300 printer called 5-1-thcolor, and the user also wants to print on one side of the sheet only, he or she should type the following command: "lpr -P5-1-thcolor -media=Tray1 -sheets=single-side mydocument.ps".

Another example would be to print the client banner: "lpr -P5-1-thcolor -job-sheets=standard mydocument.ps".

There is one limitation, however: current distributions of SLC3 and SLC4 do not have an exact driver installed for each model of the printer. Sometimes a Generic PostScript driver has been configured for the printer instead. In this case it may not be possible to select a tray or other advanced printing features. If you need to change the properties on a printer where a Generic PostScript driver has been installed, please use Windows Terminal Services instead.

Mac OS X

Users of Mac OS X have to install and maintain drivers

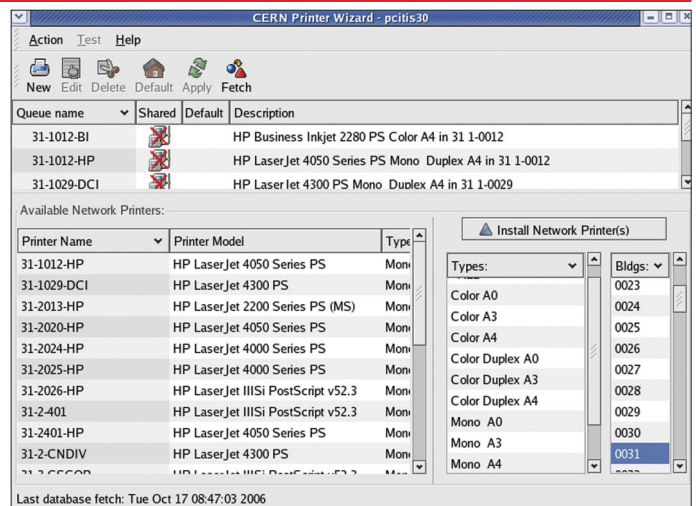


Fig. 3: The CERN Printer Wizard interface on an SLC4 computer.

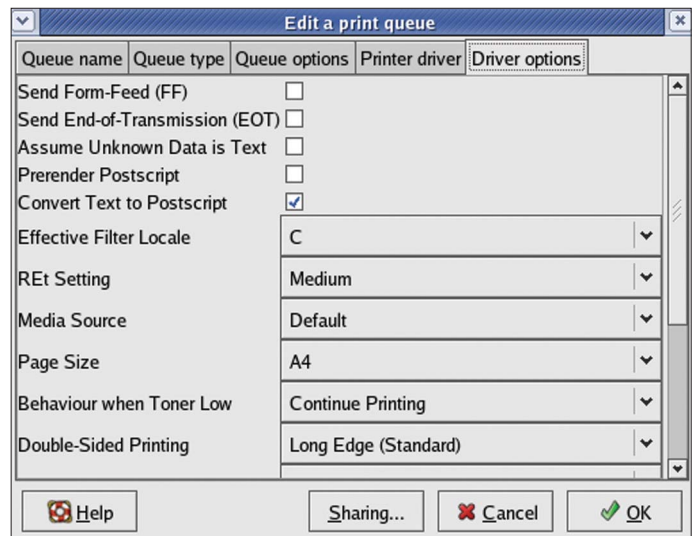


Fig. 4: Printing preferences available on an SLC4 computer.

themselves because there is no supported tool for Mac OS X such as CERN Printer Wizard on Linux SLC systems. However, drivers for most of the printers at CERN are included in the Mac OS X operating system. If you cannot find the appropriate driver, run Software Update first.

To add a new printer to your system first check its model by visiting the Printer Status page (<http://cern.ch/WinServices/Services/Printers>).

Once you know the model, open the Printer Setup Utility located in the Utilities folder. Click the Add button and select IP Printer. Now select the protocol Line Printer Daemon - LPD, and type the address "PrinterName.print.cern.ch" (PrinterName should not have the letter "B" before it). For the queue name, type "PrinterName", again without

the letter "B". Finally, select the appropriate driver from the Print Using dropdown list. Figure 5 shows how to connect to the 31-R029-HP printer.

Depending on the model, you may also be asked to specify extra features that could be installed on the printer. For instance, you could specify that the 31-R029-HP printer has a duplex unit installed, which will enable you to print on both sides of the paper.

After installation you can set a default printer by selecting a printer in the list then clicking the Make Default button. The default printer will now be displayed in bold on the list of installed printers (figure 6).

Whenever you want to change settings, you should click the Show Info button.

For detailed documentation on printing from Mac OS X,

Desktop computing

see: <http://cern.ch/it-dep/gencomputing/mac-support/PrinterSelect.htm>.

Windows Terminal Services

Printing from Windows Terminal Services is the same as printing from any NICE XP computer (see above). Under the new printing infrastructure users can install printers themselves, they do not need to ask Terminal Services administrators to do it. Each printer installed by a user is installed in his or her profile and does not affect the configurations of other users.

LXPLUS

Printing from LXPLUS is the same as command line printing with SLC3. The only difference is that all of the printers hosted by the central printing service are preconfigured on all LXPLUS nodes. Users do not have to configure printers themselves.

Visitors' Windows computers

A visitor to CERN who uses a Windows computer can also benefit from the automatic distribution of drivers. However, the built-in operating system search for the printer mechanism will not work, so we have provided a Printer Status page with search capabilities at <http://cern.ch/WinServices/Services/Printers>.

By checking the Advanced box users can search for printers by name and also by additional criteria. Once a printer has been found the user has to connect to it. Open My Computer and type “\\PrinterName.print.cern.ch\PrinterName”, where PrinterName is the name of the printer without the letter “B” before it. Now you will be asked for your credentials. Type your NICE account name and password if you have one, otherwise type “print” for the account name and “Guest@CERN” for the password. You can now change preferences and print in

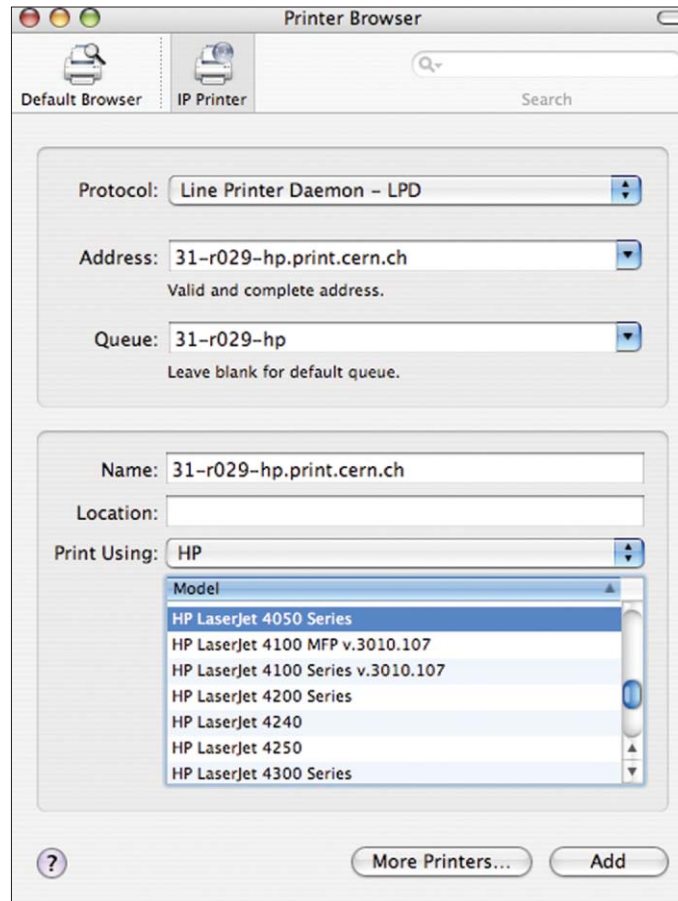


Fig. 5: Configuring a new printer for the Mac OS X computer.




Fig. 6: The list of printers configured for the Mac OS X computer.

the same way as all NICE XP users (see above).

Documentation for Windows visitors is available at <http://cern.ch/WinServices/Help/?kbid=070107#visitor>.

Visitors' Linux computers

On SLC3 and SLC4 computers a dedicated application called

CERN Printer Wizard takes care of the configuration and set-up of appropriate drivers and printers. If you have a Linux computer that is not SLC3 or SLC4, you need to install printers and drivers yourself. The first step is to identify the printer and check its model name. You should use the

Printer Status page to do this (<http://cern.ch/WinServices/Services/Printers>).

Once a printer is found you can use the information on the page to configure your system with a tool that is available on your Linux distribution (system-config-printer, printconf, YAST, CUPS web interface, etc.).

When the printer has been configured, you can print using the “lpr” command in the same way as on the SLC3 or SLC4 systems (see above).

Detailed documentation for visitors using Linux-based systems is available at <http://cern.ch/linux/documentation/printing.shtml>.

Known issues

For a list of known issues and possible solutions, consult our FAQs documentation at <http://cern.ch/WinServices/Help/?kbid=070020>.

If you have a query that is not answered on the FAQs page, please contact the helpdesk (email helpdesk@cern.ch or telephone 78888).

Further information

1. Entry point for printing service documentation: <http://cern.ch/WinServices/Help/?kbid=070001>.
2. Local CUPS documentation available on each SLC3/4 machine: <http://localhost:631/documentation.html>.
3. *Bulletin* article introducing changes to the printing service: <http://bulletin.cern.ch/eng/indexes/index.php?bullno=42/2006&base=gen&yearno=2006#Article25>.
4. *Bulletin* article on banner page policy: <http://bulletin.cern.ch/eng/indexes/index.php?bullno=42/2006&base=gen&yearno=2006#Article21>.
5. *CNL* article introducing changes to the printing service: www.cerncourier.com/articles/cnl/3/9/15/1.

Rafal Otto, IT/IS

**This is the last issue of CNL in 2006.
The editorial team would like to wish you
very enjoyable end-of-year festivities.**

Let Kerberos guard your password

Kerberos is an authentication protocol that was developed by the Massachusetts Institute of Technology (MIT) to protect network services provided by Project Athena. The protocol was named after the Greek mythological character Kerberos (or Cerberus), the three-headed guard dog of Hades.

The Kerberos protocol uses “tickets” to reduce the number of times you need to type your password each day. It has been used behind the scenes on standard CERN Linux machines for some time.

Yet it also offers advantages to other computers both on and off the CERN site, such as Macs, UNIX and Linux systems that are not centrally managed by CERN. This article explains how you can reduce the number of times you have to type your password when connecting to CERN Linux services.

What's wrong with passwords?

Passwords are an easy and attractive target for attackers. They may guess a weak password, and strong passwords can also be at risk if a password-capturing program has been installed on a compromised computer. An attacker can also listen to or “sniff” a network to intercept



Kerberos's inspiration: Cerberus, by the artist William Blake.

any passwords that are travelling unencrypted.

Once a password has been captured, an attacker can assume the user's identity. If such a password is used for multiple services, the attacker has access to all of the corresponding services (such as e-mail, accounts at other institutions and other confidential material). A compromised password can remain undetected for a long time, leaving both the user and the services vulnerable.

How can Kerberos help?

With Kerberos, users type their password once per session into their desktop computer and receive a time-limited ticket, typically for one day. To start a new session on a remote computer via SSH, instead of typing in a password, users send a Kerberos ticket. From the user's point of view access is granted “automagically” – it just works. By eliminating the need to type passwords endlessly, Kerberos makes the user's life

easier and reduces password-related security risks.

How do I configure Kerberos on my machine?

If you are running one of the supported CERN Linux systems, such as SLC4, Kerberos should already be configured.

However, if you are working on or off site and use a Mac, UNIX operating system or Linux system that is not centrally managed by CERN, then your machine may not be configured correctly for CERN. Here's how to check:

- Type “`kinit -f yourusername@CERN.CH`”.
- If you are prompted for your password and the command gives no error, then run “`klist`”. The output should include “`krbtgt/CERN.CH@CERN.CH`”.

If this is the case – congratulations, your system is already Kerberized and you should be able to access CERN machines via passwordless SSH.

However, if this is not the case, or if you receive error messages at any of these steps, take a look at <http://cern.ch/linux/documentation/kerberos-access.shtml>, where several common configurations are listed.

Kate Bradshaw, IT/DI; Jan Iven, IT/FIO

Known issues for Pine mail client users on Linux

Reconfiguring mail client

Following recent changes to CERN mail services (see p4), some Pine users have experienced problems when using the command “`uco pine -s`” to automatically reconfigure their mail client, with not all of the settings being changed. Should this happen, you should reset the configuration to CERN default. To do this:

- log on to lxplus.cern.ch;
- close Pine;
- in your home folder change the “.pinerc” file name to “.pinerc.old”;
- start Pine;
- a new .pinerc file will be created with correct settings to

read/send mails using secure IMAP/SMTP.

If you don't want to reset the Pine configuration you will need to review it manually. See <http://cern.ch/mmmsservices/Help/?kbid=191040#pine> (for advanced users).

Multiple prompts for your password

Normally Pine will ask for your password twice: once for reading (IMAP) and once for sending (SMTP). Additional prompts often indicate that Pine is using both “imap.cern.ch” and “mmm.cern.ch” aliases to access your mailbox. Make sure that wherever the path to your mailbox is mentioned in Pine

configuration, it always starts with the same server, preferably “imap.cern.ch”.

Address book

Pine does not support authenticated LDAP access to an address book server. Therefore, Pine users can only query the address book from inside CERN, where anonymous LDAP is still allowed. If you want to use Pine from outside CERN with the address book feature available, you can connect remotely to lxplus.cern.ch then start your Pine client from that machine.

Bounce

The Bounce command cannot

work with authenticated SMTP. Users should now use the Forward command. It is no longer possible to bounce mails.

Other options for reading e-mails

Linux distributions at CERN are all shipped with graphical and fully featured mail clients, and seamlessly support secure protocols. Supported products are Thunderbird and SeaMonkey, formerly known as Mozilla – as explained in the last issue of *CNL* at www.cerncourier.com/articles/cnl/3/9/17/1. For more details on SeaMonkey look at: <http://cern.ch/mmmsservices/Help/?fdid=18>.

The User Support and Mail teams

Users can gain access to AIS with NICE password

As announced in the June–August issue of *CNL*, it is now possible to use your NICE username and password to log on to AIS (see www.cerncourier.com/articles/cnl/3/6/14/1). The procedure is now fully operational and users can reset the AIS password themselves so that their NICE password can be used for AIS authentication.

We strongly recommend that CERN users who have a NICE account do this, to reduce the number of passwords they need to remember. This can be done easily by selecting the Change Password option that is displayed at the bottom of the AIS log-on page (<https://aislogin.cern.ch>). Provide the Old Password, click on the button Use Nice Password, then finally click on Submit (figures 1 and 2).

The procedure enables the NICE password to be used for AIS authentication, but not the other way round. Please look at <http://cern.ch/security/passwords> to select a good password that will be accepted by the NICE system.

EDH authorization password

A second password is still needed to sign documents in Electronic Document Handling (EDH). Known as the “EDH authorization password” or “AIS signature password”, this is your digital signature, which replaces the written signature that was used in the past.

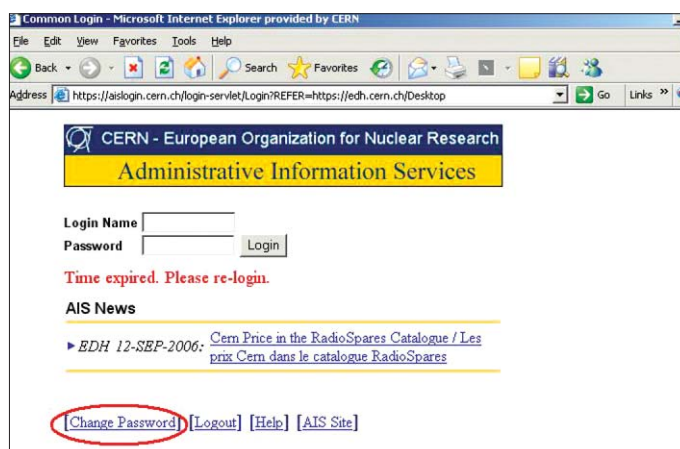


Fig. 1: Select the Change Password option on the AIS login window.



Fig. 2: Reset the AIS password to your NICE password.

If you want to reset your EDH authorization password you must first access the AIS system (<https://aislogin.cern.ch>) with your AIS log-on password (preferably reset to use the NICE password), and then go to <https://edh.cern.ch/Desktop/settings.jsp> (figure 3). There is



Fig. 3: Your EDH authorization password should be different.

no link between the AIS log-on password and the EDH authorization password, the two are unconnected. For security reasons we recommend that you continue to use different passwords for logging on to AIS login and your EDH signature.

The User Support team

Calendar

November

11–17 **SC06: International Conference for High Performance Computing, Networking, Storage and Analysis**
Tampa, Florida, USA
<http://sc06.supercomputing.org>

21–24 **CIC 2006: 15th International Conference on Computing**
Mexico City, Mexico
<http://magno-congreso.cic.ipn.mx/CIC-2006>

22–24 **IST 2006: Information Society Technologies Conference and Exhibition**
Helsinki, Finland
http://europa.eu.int/information_society/activities/istevent/index_en.htm

December

3–6 **LISA'06: 20th Large Installation System Administration Conference**
Washington, DC, USA
www.usenix.org/events/lisa06

4–6 **e-Science 2006**
Amsterdam, Netherlands
www.gridbus.org/escience

2007

February
13–15 **PDCN 2007: IASTED International Conference on Parallel and Distributed Computing and Networks**
Innsbruck, Austria
www.iasted.org/conferences/home-551.html

March

3–6 **WEBIST 2007: 3rd International Conference on Web Information Systems and Technologies**
Barcelona, Spain
www.webist.org

April

23–27 **ACAT07: Advanced Computing and Analysis Techniques 2007**
Amsterdam, Netherlands
www.nikhef.nl/acat07

23–25 **Spring 2007 Internet2 Member Meeting**
Arlington, Virginia, USA
<http://events.internet2.edu/2007/spring-mm>

The deadline for submissions to the next issue of CNL is 15 December

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>