


CERN COMPUTER NEWSLETTER



CERN - IT Division
CH-1211 Geneva 23
SWITZERLAND

July – December 2002

CERN-CNLI-2002-003
Vol. XXXVII, issue no 3

Contents

If you need Help	i
Editorial	iv
1 Announcements	1
1.1 Closure of Off-Site Telnet	1
1.2 End of NICE 95 and NICE NT Services	1
1.3 Computing Helpdesk – a New Strategy	2
2 Physics Computing	4
2.1 DataGrid Moves to Production	4
2.2 Data Services Group Activities and Plans	5
2.3 AFS Project Space Administration - New Project Quota Structure	6
3 Desktop Computing	7
3.1 CERN Linux 7.3.1 Certified	7
3.2 Solaris 8 Certified	7
3.3 Tru64 Unix (Dunix) Support at CERN	8
4 Internet Services and Network	9
4.1 A New Infrastructure for the CERN Electronic Messaging Service	9
4.2 Receiving Unwanted Email?	9
4.3 Batch Nodes on Non-Routable Network	10
5 Scientific Applications and Software Engineering	11
5.1 Anaphe – OO Software for HEP in C++ and Python	11
5.2 2002 CERNLIB Release	13
5.3 EXtreme Programming (XP): Pragmatism & Quality	14
6 The Learning Zone	15
6.1 Questions and Answers from the Computing Helpdesk	15
7 User Documentation	20
7.1 News from the User Services Bookshop	20
7.2 User Services Book Catalogue	20

Merry Christmas!

Happy 2003!

Editorial Information

Editorial Board

Chief Editor	Nicole Crémel (CNL.Editor@cern.ch)
Associate Editors	Charles Curran Ludwig Pregernig Frédéric Hemmer
Technical advisor (printed edition)	Michel Goossens

The opinions expressed in this newsletter are those of the contributors and are not necessarily those of the CERN Management. The Editors reserve the right to edit or omit articles, or hold-over copy due to lack of space.

General Information

The CNL Web home page is at URL: <http://cern.ch/cnl>

The policy is to have between two and four CNL editions per year.

Issue number	Deadline
CERN-CNL-2003-001	Friday 14th March 2003

Information for contributors

The CNL is processed using XML technology and templates. Articles marked up in HTML are preferred, but plain text is also acceptable. A Web form to be used for article submission (contributions) is provided at the CNL home page as a link “Article Submission form”.

Distribution

- An electronic mail is sent to the “CNL Announcement Distribution List” as soon as a new issue of the CNL is available on the Web.
- A paper version is sent to the users who have subscribed to the “CNL Paper Version Registration” list.

All *subscriptions to our distribution lists* can be made through the “Help&Info” link on the CNL home page.

You can also cancel your subscription in either case by filling in the appropriate form.

The compressed *PostScript file* containing the complete printable version of this CNL is available via the Web at the URL <http://cern.ch/ref/cnl-2002-003/main.ps.gz>. It is possible that, in order to optimize typesetting, the PostScript printable version differs slightly from the master HTML version. In any case the HTML version on the Web must be considered as the reference.

If you need Help

FRONT LINE Services

	Phone	Location	E-mail Address
Computing Help Desk Opening hours: 8:30-17:30 (Monday-Friday)	78888		Helpdesk@cern.ch
Administrative Applications (EDH, BHT, etc.) Opening hours: 8.30-12:30 and 13:30-17.30	79933	5/R-021	Ais.Support@cern.ch
Central Computer Operators 24/24 hours on call	75011	513/R-049	Computer.Operations@cern.ch
General Network (Operational Problems)	74927		Network.Support@cern.ch

Contacts at CERN

The “Computing Help Desk” is the call center for help on IT services related issues. Following is a non-exhaustive list of services that are provided by the IT division, and the people responsible for it.

Please note:

- In the Web version of this page the "Service Definition" can be a link to the Welcome page of this service: we invite you to read carefully all the information given in this page before submitting a question/problem.
- In addition you might find useful to search in the “Questions&Answers” (<http://cern.ch/consult/qas>) database for possible answers to your query.
- Many generic e-mail addresses for support lines are just an automatic re-direction to the Computing Helpdesk (this is the case, for instance, for Mail.Support or Www.Support). Those addresses have been kept mainly for backward compatibility, and also to help the Helpdesk to escalate the query towards the right specialist whenever this is needed.

USER SERVICES

Service Definition	Name	E-mail Address
User Relations and Service Manager	L.Pregernig	Ludwig.Pregernig@cern.ch
Help Desk Managers	R.Woolnough	User.Relations@cern.ch
	N. Crémel	User.Relations@cern.ch
User Registration		User.Registration@cern.ch
Computing Bookshop	J.Megies	Bookshop@cern.ch

CENTRAL SERVICES

Public Services (PLUS batch&interactive services and CSF/PCSF Simulation farms)

Service Definition	Name	E-mail Address
CORE Operations	M.Vergari	Mario.Vergari@cern.ch
Distributed File Services (incl. AFS)		Afs.Support@cern.ch
NAP and Parallel Applications	E.Mcintosh	Eric.Mcintosh@cern.ch
Stager and Tapes Support		Tape.Support@cern.ch
Central Data Recording (CDR)		Cdr.Support@cern.ch
CASTOR HSM software		Castor.Support@cern.ch
Print Service	J.L.Vosdey	Printer.Support@cern.ch

Support for Experiments

You can see the complete list of services and contacts at URL: <http://wwwinfo.cern.ch/pdp/is/main/experiments.html>

DESKTOP COMPUTING

Service Definition	Name	E-mail Address
NICE (CERN Windows Services)		Nice.Support@cern.ch
PC Desktop (Repairs and Sales)	C.Ball	PC.Admin@cern.ch
MACintosh Support	F. Ovet	Mac.Support@cern.ch
MACintosh Sales	W.Hug	Waltraud.Hug@cern.ch
X Terminals Support		Xterminal.Support@cern.ch
ASIS Support		Asis.Support@cern.ch

Operating System Support:

Service Definition	Name	E-mail Address
HP 700		Hp.Support@cern.ch
SUN		Sun.Support@cern.ch
Solaris		Solaris.Support@cern.ch
RS/6000		Aix.Support@cern.ch
Digital Unix		Dunix.Support@cern.ch
Linux		Linux.Support@cern.ch

SCIENTIFIC APPLICATIONS & SOFTWARE ENGINEERING

Service Definition	Name	E-mail Address
CERN Program Library	L.Besson	Heplib.Support@cern.ch
PAW Support	O.Couet	Paw.Support@cern.ch
Anaphe		Heplib.Support@cern.ch
Oracle Databases	N.Segura	Oracle.Support@cern.ch
Objectivity DB Service		Objectivity.Support@cern.ch
Software Development Tools		SDT.Support@cern.ch
Adobe FrameMaker		Docsys.Support@cern.ch
IT Controls		ITControls.Support@cern.ch

COMPUTING FOR ENGINEERING

These activities are now in the Product Support (PS) group.

Index of tools: http://cern.ch/product-support/support_index.html

COMMUNICATIONS AND NETWORKS

See the “Communications Systems Group” web pages, at URL <http://cern.ch/it-div-cs/public/help/> and <http://cern.ch/it-div-cs/>

INTERNET SERVICES

Service Definition	Name	E-mail Address
Web Services		Www.Support@cern.ch
Mail Support		Mail.Support@cern.ch
MMM (Exchange 2000)		Mail.Support@cern.ch

In addition to this generic address, the normal support lines are listed at URL:
<http://cern.ch/mailservices/Support/>

COMPUTER SECURITY

Computer Security Information	Computer.Security@cern.ch
Security Alert Reports	cert@cern.ch

Other Newsletters

Administrative Information Services: CERN/AS AIS Newsletter (<http://ais.cern.ch/newsletter>)
Contact: AIS.Newsletter@cern.ch

Some Useful Web Pages

IT Home Page	http://cern.ch/it-div
TvScreen (Systems Status)	http://tvscreen.cern.ch
Network Connection Request Form	http://network.cern.ch/register
Remote/Mobile Computing	http://cern.ch/it-div-cs/public/services/mobile/
User-Oriented Catalogue	http://cern.ch/consult
User-Oriented Service Directory (*)	http://cern.ch/services
FrameMaker Support	http://cern.ch/docsys/framemaker/
L ^A T _E X at CERN	http://wwwinfo.cern.ch/asdoc/textproc.html
NICE (CERN Windows Services)	http://cern.ch/winservices/
PC Hardware	http://pcwww.cern.ch/
CASTOR Query Report	http://wwwinfo.cern.ch/htbin/pdp/castor/castor_report.pl
GEANT4	http://cern.ch/geant4/
Software Development Tools	http://cern.ch/sdt
Joint Controls Project (LHC Experiments)	http://itcowww.cern.ch/ITCO/jcop/
PVSS Service	http://cern.ch/itcobe/Services/Pvss/
Mail Services	http://cern.ch/mailservices/
MMM Services (Exchange)	http://cern.ch/MMM/
Web Services	http://cern.ch/WebServices
Computer Security Information	http://cern.ch/security

(*) Useful service entry points at this page include the Computing Helpdesk, User Registration, Computing Bookshop, Mail, Netscape, Listbox, etc.

Seminars and Presentations

Designation	Usual Location	Organizer(s)	E-mail Address
Computing Colloquia	Auditorium / bldg 500 see http://cern.ch/computing-colloquia	J. Shiers, IT	Jamie.Shiers@cern.ch
Computing Seminars	IT Auditorium 31/3-005 see http://wwwinfo.cern.ch/seminars/	J. Blake, IT	computing.seminar@cern.ch
Technical Presentations	IT Auditorium 31/3-005	S. Cannon, IT	stanley.cannon@cern.ch

Computer Resources Allocation and Divisional Contacts

Computer Resources Allocation Committee (COCOTIME):

Chairman:	H. Hoffmann/DG	60/5-006	75458	Hans.Falk.Hoffmann@cern.ch
Secretary:	A. Morsch/EP	160/R-018	78617	Andreas.Morsch@cern.ch

Divisional Representatives for Computing

For questions about new registration and resource allocation please ask the contact person for your division.

Div.	Repres.	Div.	Repres.	Div.	Repres.	Div.	Repres.	Div.	Repres.
AS	Mats Møller	EP	A. Bjørkebo	EST	C. Andrews	ETT	M. Draper	FI	C. Saitta
HR	M. Buttner	IT		PS	A. Riche	SL	M. Bornand	SPL	F. Costa
ST	E. Sanchez-Corral	TH	E. Gianolio	TIS	G.R. Stevenson				

Divisional Representatives for Desktop Forum

See the URL <http://cern.ch/mgt-desktop-forum/Members.htm>

Divisional Representatives for Networking

For questions about new installations, testing, local coordination, etc., please ask the contact person for your division (or, if you are an experimental physicist, the networking contact person within your experimental group).

The complete and up-to-date list of contacts is given in the “Communications Systems Group” Help page, at the URL:
<http://cern.ch/it-div-cs/public/help/>.

Editorial

Dear CNL readers,

For the originally planned “Autumn” issue of the Computer Newsletter, we received fewer contributions than expected. Therefore, we decided to publish one combined (“Autumn” and “Year End”) issue, and it is this CNL which is now in front of you. It contains several important announcements and many interesting articles:

- several announcements and changes related to telnet closure, end of Windows 95 and NT at CERN, and a new model for the Computing Helpdesk,
- a detailed presentation of the European DataGrid project, and its present status,
- a description of the new infrastructure that will be provided soon by the Mail and Web services,
- some advice on recognising SPAM (unsolicited emails) and what to do about it,
- a description of the Anaphe project, to provide an Object Oriented software environment for HEP experiments.

I hope you will enjoy reading this new issue, and, at this time of the year, I wish you a ...



... and all the best for a happy 2003!

Nicole Crémel, CNL Editor, IT/User Services



1. Announcements

This chapter contains announcements concrete in time and/or of informative nature from the division to the User Community.

1.1 Closure of Off-Site Telnet

Denise Heagerty, CERN Computer Security Officer

Abstract

Off-site telnet access to CERN will be blocked on Tuesday 28th January 2003.

To reduce the number of regular break-ins on CERN machines due to passwords exposed on the network in clear text, **Off-site telnet access to CERN will be blocked** in the CERN firewall from

Tuesday 28th January 2003.

Users are recommended to install and use SSH (Secure Shell) as an alternative to both the telnet and ftp commands as soon as possible for off and on site access. Further details are at the URL <http://cern.ch/security/ssh>.

In addition, please note the following:

- FTP also exposes passwords in clear text and is planned to be closed at a later date when alternatives are available for all services (e.g., NICE file servers and central web servers still require ftp). SSH (the commands scp or sftp) provides similar file transfer functionality to ftp and is a recommended alternative where this is possible. It is more secure and with compression can also be more efficient.
- You are recommended to change passwords which could have been exposed by telnet and ftp as soon as you convert to SSH.
- ACB (<http://network.cern.ch/acb>) and VPN (<http://cern.ch/vpn>) access are considered as if *on site* and are not affected by this telnet block. You are however strongly recommended to use SSH to access CERN systems for all connections, whether on or off site.



1.2 End of NICE 95 and NICE NT Services

Christian Isnard, IT/IS

Abstract

NICE 95 and NICE NT services will be frozen at the end of January 2003 and will stop on 1st April 2003.

You are concerned by this article if you still use a computer running NICE 95 or NICE NT.

As recommended by the Desktop Forum, NICE 95 and NT services will be stopped according to the following schedule:

31st January 2003.

NICE 95/NT will be frozen. Applications will still be able to run, but the helpdesk will not address any NICE 95/NT problems any more. It will not be possible to reinstall NICE 95/NT any more. Disk images of the *original* Windows 95 and Windows NT CDs are available on the network, but it will be up to the user to create those CDs, reinstall the machine(s), as well as to locate and install:

- the required applications,
- the device drivers for special hardware,
- the necessary security patches,
- anti-virus software: Operational Circular N°5 (http://cern.ch/ComputingRules/OC5_english.pdf) and the CERN Computing Rules (<http://cern.ch/>)

ComputingRules) require that Windows PCs are regularly checked for viruses.

The original Windows 95 and Windows NT CD images can be obtained from the URL <http://cern.ch/win/services/installation/CDImages>. You should give your NICE username and password to access this page.

Useful security information and patches can be found at the URLs:

- <http://www.microsoft.com/networkstation/downloads>
- <http://www.microsoft.com/windows95/downloads>
- <http://cern.ch/security>

On April 1st 2003:

The NICE 95/NT servers will be stopped. Applications will not be able to run any more unless they have been configured for standalone use before 1st April (only possible with Office 97, Netscape 4.x, Exceed 6, Winvn, PC Pine 3.x, Publisher, Filemaker, Framemaker, Mathematica via the “*PC Configuration*” tab of the “*CERN Settings*” program available in the control panel).

It will thereafter be the responsibility of the user to keep his/her system and applications secure (see links to security information in the previous paragraph). A system that is not kept secure can be required to disconnect from the network in order to protect the CERN site.

The support for generic services (like access to home and divisional directories, electronic mail, printing) will be limited to NICE 2000.

You are therefore advised to take action urgently to ensure that the remaining NICE 95/NT computers that you may use

are migrated to NICE 2000, replaced or definitively switched off.

In case of doubts, or if you have any questions — in particular concerning any specialised application or hardware that you may have problems to migrate to NICE 2000 — please get in touch with your Desktop Forum representative (for more details see the URL

<http://cern.ch/mgt-desktop-forum/Members.htm>.

Thanks for your cooperation from the NICE team.



1.3 Computing Helpdesk – a New Strategy

Nicole Crémel and Roger Woolnough, IT/User Services

Abstract

This is a follow-up to the article “Move and Upgrade of the Computing Helpdesk” published in the last CNL, to explain the new model that has been put in place since then.

Move and Upgrade - Next steps

The changes that have been described in the previous CNL (see “Move and Upgrade of the Computing Helpdesk” at the URL <http://ref.cern.ch/CERN/CNL/2002/002/helpdesk-upgrade/>) and made for a trial period, have been put in place definitively since June 2002.

Changing the Helpdesk from a “dispatch” to a “resolution” centre and to ensure consistent quality replies, we have decided to re-centralise the Helpdesk in one place. This means that the satellite desk in Building 52 has been definitively closed on September 30th 2002. In compensation, the improved Helpdesk in building 513 includes more staff and is better equipped to deal with user questions.

We appreciate the replies from users to our “request for feedback”. On the average, about 10% of users give us feedback: 53% of the replies are “Thank You” 44% requests for solution updates, and 3% are complaints. Please continue this feedback — either replying to the message sent by the system, or sending a mail to User.Relations@cern.ch (IT User Services Group) — - this helps us to improve the service for the future.

The IT Computing Helpdesk operates Monday to Friday, from 08.30 to 17.30 nonstop. To contact the service users must either send an email (to Helpdesk@cern.ch) or phone (78888). When sending a message, please include as much detail as possible about your problem; it helps to find a solution more quickly. Please remember that the Helpdesk’s role is to find solutions for problems with established services from the IT division. Solving other problems is outside the scope of the Helpdesk, but we will inform service providers on issues brought up in this context.

A New Model for more Quality

In the past, the quality of the answers provided by the Helpdesk or service providers has never been rigidly controlled. This has now been changed. The Helpdesk contract staff (outsourced) are now backed up with more skilled staff to improve technical knowledge. Since September 2002, two important changes have been made:

1. Firstly another outsourced specific contract has been defined for “mail and web support” which is an important fraction of the helpdesk calls. This person sits in the same room as the Helpdesk and can help immediately with difficult questions raised by phone.
2. Secondly, one member of the User Services group in IT is working closely with the helpdesk team, in order to ensure a better service and to ensure some quality control. With the presence of a CERN staff supervising the outsourced team (we call him/her the “Manager on Duty”) more questions are treated directly at the source, without further escalation.

About six members of the User Services group are currently sharing the “Manager on Duty” (hereafter called MoD) task on a shift rotation scheme, but more people in the IT division will participate in the near future.

A criticism of the Helpdesk, and of the Problem Tracking System, has often been its “anonymous” behaviour: users did not know whom to contact if their case was incorrectly treated. Now, for all the e-mails that go through the helpdesk, in the acknowledgement message users can see the name of the MoD who is formally “responsible” for the ticket that has been created in the Problem Management tool. “Formally responsible” does not mean that the MoD will provide the reply and solution, but, in case there is a problem with that ticket (delay is too long, the proposed solution was not good, etc.), users can address any feedback or complaint to this person, who will then ensure a proper follow-up.

The same Problem Tracking System (Action Request System, Remedy) is also used by other services. If the acknowledgement message does not contain a name, but the e-mail

User.Relations@cern.ch instead, it means that your mail did not go through the Helpdesk, but directly to the service concerned (for instance AFS.Support, or Linux.Support). In this case the MoD and the Helpdesk have not been involved in the process for problem resolution. However if some follow-up is needed, it remains the work of the User Services group via the user relations team.

Role and tasks of the Helpdesk Duty Manager (or MoD)

The first role of the MoD will be to ensure some quality control. Here are the practical actions, already implemented, in order to achieve this:

- Sitting in the helpdesk office, the MoD can help the staff answering phone calls with difficult questions, or where the helpdesk staff have problems understanding or explaining. If necessary users can also ask to talk to the MoD directly.
- The MoD is systematically receiving copies of all the “Possible solutions” that are sent to the users, and copies of the mail replies sent by the users. When needed the MoD will ensure the follow-up (for instance if the possible solution requires an update).
- The MoD coordinates the resolution process with service providers. Reminder mails for cases which are not followed according to the divisional service-level agreements are automatically sent both to the service provider who has been assigned the case, and to the MoD attached to this case.
- The MoD acts as the link between users and service providers. He/she will contact service providers in case of problems that might be identified by the Helpdesk, the operators or other sources, in order to co-ordinate user information (news, zephyr messages

or other user contacts).

Service providers are requested to contact the MoD as soon as possible in case of interrupts so that users may be contacted and messages co-ordinated.

Objectives for the future are more ambitious and we expect that, in a second stage, the duty manager should:

- Prepare a report of problems that need to be tackled by divisional management. This report should highlight anything involved in improving the service quality from the Division. This report should be coupled with statistical information.
- Maintain a service status board and web page for user information. The board should contain information on actual problems and forthcoming changes. The latter should at least be mirrored on the web page and tvscreen tool.
- Co-ordinate a “knowledge base” where users could find answers, without calling the helpdesk. Until now the XFIND tool (<http://consult.cern.ch/xfind>) for the domain “Public – QAs” has been used as such a database. A project is under way to review all Q&As and, using the appropriate technology, to come up with an efficient tool for user self-help. The MoD should ensure that when QAs are missing a new one is created, or, check that old QAs are still valid or if necessary re-edited. This work is done in collaboration with the service providers and is an on-going task. Again, users who find discrepancies are requested to let us know through the User.Relations email.
- Finally, after more experience at the helpdesk, the duty manager should be able to suggest ways whereby the Helpdesk receives fewer calls. Areas that have been identified are automating password resets and space provisioning (“quota”).



2. Physics Computing

Functional description of all the Central Public Services, specifically oriented to use by the Research sector, either for interactive (central and CPU intensive) or batch processing. The main physics-computing services consist of providing computer farms, central data recording facilities and central data storage facilities.

2.1 DataGrid Moves to Production

DataGrid Project Office, CERN

Abstract

The DataGrid project is moving towards the production of a stable resource that scientists all over Europe will be able to plug into and use.

When the European DataGrid (EDG) project started, almost two years ago, one of its main aims was to go beyond the R&D phase and demonstrate a "production quality" computing Grid. Such a Grid would give scientists unprecedented computing power to tackle major challenges, such as modelling climate change or analysing genomic data, using a distributed network of computers involving major computer centres all over Europe. Production quality means not just a proof of principle, but the deployment of a stable resource that European scientists would be able to plug into and use on a regular basis.

In the last two years, those ambitious goals have come a lot closer to reality.

Already in March this year, during the first EU Project review, a demonstration was made of the first release of DataGrid middleware—the software that makes a Grid of computers work together seamlessly. Jobs were submitted to the DataGrid Testbed1, showing the capability of the grid to distribute computational tasks to the most appropriate resources, by matching the set of requirements specified by the user to the characteristics and status of the available resources (CPU power, memory size, CPU load).

During the CERN school of Computing (15-28 Sept 2002, Vico Equense (Italy)), 80 students have been exercising the functionality of the software, submitting jobs during the hands-on session dedicated to the Grid.

A new release of the Grid middleware was officially made on November 11th, with greatly improved support for large file transfers, better tracking of applications as they run on the Grid and a more stable information system. Several key features for production have been added, such as simplified access to mass storage systems, an easier software installation mechanism and user friendly job submission facilities.

The software is currently being installed on hundreds of computers that make up the EDG production testbed. Initially limited to five European sites, the testbed has recently been enlarged to approximately 20 sites across Europe, including Poland, the Czech Republic, Spain, Portugal, Germany and the Nordic countries.

In the months to come, the focus and resources of the project will be concentrated on supporting the LHC experiments that are preparing for the future data taking and analysis by running the so-called Data Challenges, a set of computing tasks concerning software simulation of physics processes. The status of this work will be reported at the second EU project review scheduled for the 4th-5th February 2003.

An ATLAS-EDG Task Force was put together in August 2002 with the aim to assess the usability of the EDG testbed for the immediate production tasks. The first phase of "Data Challenge 1" has been already performed using various Grid tools and results have been presented during the September 19 Task Force meeting. The second phase, starting in October 2002, will use Grid tools as much as possible.

CMS is starting now similar usability evaluation, and the same is happening in the other LHC experiments and will be supported by the DataGrid experts. The hope is of course to contribute middleware technology and testbed infrastructure to the overall LHC Computing Grid project.

In the mean time, Earth Observation scientists are busy developing the EO Grid infrastructure, and deploying a first Grid-based prototype for collaborative processing and validation of atmospheric ozone levels measured from ESA satellites. Biomedical applications, which were already successfully demonstrated at the time of the first project review, continue to be developed and enhanced.

A practical introduction to grids and to the EDG software is available for anybody who participates in the "DataGrid tutorial", a *travelling* event that, after having been featured already twice at CERN in October, will move to several institutions in the project members states. The tutorials are organized over two full consecutive days and include both a theoretical introduction on general grid concepts, user interface commands, grid components, etc., and hands-on exercises covering user interface operation with job submission, aspects of data management and job monitoring. No prior knowledge is required for the participation.

For more information on the DataGRID project see:

<http://www.eu-datagrid.org>

For the programme and schedule of the tutorial:

<http://eu-datagrid.web.cern.ch/eu-datagrid/Tutorial/tutorial.htm>



2.2 Data Services Group Activities and Plans

Harry Renshall, IT/Data Services

Abstract

This is an update on the IT Data Services group activities and plans and covers AFS, backup, CASTOR and tape services.

The Data Services group is currently responsible for four major areas. These include the home and project directories infrastructure based on the OpenAFS product and the backup services for computer centre and departmental servers based on the TSM and Legato products. We are responsible for the managed storage software and services of the CERN Advanced Storage Manager (CASTOR) (which includes backwards compatibility with the SHIFT tape access software) and for the magnetic tape drive, robotics and disk server infrastructure that underpins all of these. Changes are planned or have recently happened in all of these areas and it is timely to report on them.

AFS Services

During the first quarter of this year most computer centre AFS client machines were changed from IBM to OpenAFS and this improved the stability of the client and fixed some of the problems where a client would lock up completely any access to AFS. We were still left with a serious bug in handling large directories that were frequently changed from multiple clients (e.g., hundreds of concurrent batch jobs) and this was found to be on the server side. This was fixed in a new release of the OpenAFS server code and we installed a spare machine with this and moved over the busiest directories. Since then we have been systematically migrating our twenty-five servers and have now completed this work. Each server migration took several weeks as live user data was slowly but transparently drained from it to an already migrated server.

Backup Services

Our major activity in this area is planning to reduce from the current three systems we use (AFS internal, IBM Tivoli Storage Manager and Legato) to a single system next year. This will probably be, though not necessarily, one of the existing systems. This will save both manpower and licence fees and is made possible by the TSM and Legato products converging in performance and functionality. In December we will present a proposal for implementation next year. This will be largely transparent to the ordinary central computing users.

Managed Storage

The CASTOR system now handles the bulk of the physics data of CERN with a total volume stored on tape of more than 1 Petabyte and over 7 million files. The software and services using it have reached a good level of stability and performance recording, for example, from 2 to 4 TB of data

per day for the COMPASS experiment during their data taking this year. An important test will be the Alice Data Challenge which, during November/December, intends to store data to tape at 200 MB/sec for a continuous week.

A user meeting was held in March 2002 and a prioritised program of work was built based on user and operational requirements. User priorities were to support the LHC Data Challenges and improve reliability and we are introducing large file support, tape re-packing, tape drive fair shares allocation and improved monitoring and statistics. The main items can be seen under the link “Development Plans for the rest of 2002” on the CASTOR home page at the URL:

<http://it-div-ds.web.cern.ch/it-div-ds/HSM/CASTOR/>

An important item is to provide a GridFTP interface to CASTOR to run on a second `wacdr.cern.ch` gateway server and this should be ready in November.

Tapes and Tape Robots

In January of this year we completed the symmetric splitting of our STK tape robots and drives into two physically separate complexes, one on the ground floor of building 513 and the other in a new building 613 some 200 metres away. Each complex consists of five silos with each silo having 5500 storage slots and each complex has 14 of our mainstream tape drive, the STK 9940A. These drives move data to and from tape at 10 MB/sec and store 60 Gigabytes per cartridge. STK have now announced the 9940B model drive which moves data at 30 MB/sec and stores 200 GB onto the same cartridge as used by the model A drive. The LCG project ordered 20 of these drives to be used for the Alice data challenge and they were delivered to CERN on 13 November. We have already performed an extensive field test and found they perform reliably and as specified. Next year we hope to upgrade all our 9940A drives to B models and reduce our media storage costs from above FS 2 per GB on tape to below FS 1. The model B drive can read a tape written on the model A but cannot write at the lower density. The upgrades are planned to be ready for the next SPS run from May 2003.

We announced last year the ending of the STK Redwood helical scan magnetic tape service at the end of 2002 (these are tapes with visual identifiers beginning with the letter Y) and we are now in the last stages of this operation. We have already copied most required Redwood tapes into the CASTOR system. All Redwood tapes except for those of NA48 have been locked in our Tape Management System (TMS). The NA48 experiment have 7000 Redwoods of 50 GB each, representing several years of raw data. We had been waiting until we had some of the higher density 9940B drives to copy them but STK kindly allowed us to keep some of the field test drives while waiting for our order to arrive and we have been using those. We are now ejecting the already copied or unwanted Redwood cartridges and putting them in

storage. There will be no maintenance possible on Redwood drives from the end of this year and given their short head life we cannot guarantee to be able to re-read any Redwoods next year.

A third major activity will be to provide the resources for the Data Base group to migrate 2001 and 2002 COMPASS data (over 3000 cartridges) from Objectivity format to Date format and create an associated metadata database of a few TB using Oracle. As with the Redwood copies for NA48 this will be done onto the high density 9940B drives so can only start when the Alice Data Challenge has liberated enough drives. We estimate the copying will take 2 to 3 months. Note that the COMPASS and NA48 experiments have had 12 9940A tape drives dedicated to them during this year's SPS run and we plan for eight of those to be dedicated to be the input drives for this COMPASS migration exercise. This temporary dedication is not expected to have a serious effect on other users.

A fourth major activity will be to relocate the tape robots in the building 513 computer room to their final position in the former tape vault of building 513. The vault was emptied at the end of last year and major refurbishment work has been carried out to equip it to take up to eight tape robots and thousands of PCs. This is part of the long term plan to also equip the ground floor computer room in building 513

to be ready for the LHC computing equipment. This will be done in three phases and relocation of the tape robots is part of the first phase. This work must also be completed well before the start of the SPS next year and we plan to do it in January and February. There will be some disruption to tape mounting but we plan to minimise this by emptying complete silos, made possible by ejecting all Redwood cartridges, and moving them one at a time.

Finally, as well as the tape robots there are some eighty disk servers (out of a total of over 200), many of which also run stagers, that have to be moved early next year. These were the first ones installed and are spread over nearly all Cern experiments. Among them are the primary CASTOR name server machine and four new name-servers scheduled for each LHC experiment next year. The servers are mixed over many experiments and are spread over three network services. Their network infrastructure must be dismantled and moved in synchronisation. It is not physically possible to cleanly move one experiment at a time so we propose to shut down all access to the CASTOR system for the time required to make the moves which is estimated to be three days. We are proposing the dates of Monday/Tuesday/Wednesday 13/14/15 January 2003 as the workload is traditionally lower in early January. We will, of course, try and minimise the actual time and will publish more details later.



2.3 AFS Project Space Administration - New Project Quota Structure

Rainer Többicke, IT/Data Services

Abstract

The command `afs_admin` to manage AFS disk space.

With the new `afs_admin` command introduced last summer a revised structure has been put in place allowing for refined control of AFS project space allocation. `afs_admin` is the tool used by AFS project space administrators to manage experiment/project disk space, as well as scratch/work space given out to individual users from the experiments quota.

Formerly, only two distinct categories of AFS project space ("backed-up" and "non-backed-up") were available. They were controlled by a single "project quota" for the experiment/project limiting the aggregate quota of AFS volumes in the project. AFS volumes in the "backed-up" category would typically start with `'p.projectname...'` and in the "non-backed-up" category with `'q.projectname...'`.

In the new structure:

1. Quota for each category can (would normally) be specified individually.
2. A new category "user-scratch" has been added: meant as scratch work space for individual users, usually implemented on cheap disks. In case the disks or the servers holding them fail, the data may be lost, as there are no backups. This space is therefore unsuitable for data which cannot be re-created or recovered

by other means. Scratch volumes are created using the `afs_admin create_scratch` command. Volumes in this category start with `'s.projectname...'`, and will have a directory name of `~userid/scratchN`, `N` being 0, 1, 2, ...

3. The previous categories remain as before: `'q.projectname...'` volumes are for data which will not be backed up and should therefore be re-creatable by other means, `'p.projectname...'` and `'u.projectname.userid.X'` volumes (accessible under `~userid/wN`, with `N=0,1,2,...`) are backed-up daily to tape, back-ups are kept for about a year at increasingly spaced intervals. Data in these two categories are placed on highly reliable servers and disks.

`afs_admin list_project` and `afs_admin list_quota` are commands that now list the project's AFS volumes or its "project quota" under the different categories separately—provided distinct values have been assigned—with headings `'projectname_q'` or `'projectname_s'`, or even `'projectname_p'`. Additionally, it is now possible to restrict the output of the command `afs_admin list_project` to volumes of a certain category by appending a `'_p'`, `'_q'` or `'_s'` to the project name, e.g., `afs_admin list_project atlas_q`

3. Desktop Computing

Description and news concerning the CERN desktop environment centrally provided on UNIX and NICE/PCs (Windows-95, -NT or -2000), and providing the “general-purpose computing” environment. It consists of the hardware, software strategy, and services for the basic applications, which one expects to be generally available on desktops. This environment presents the basic computing infrastructure of the laboratory.

3.1 CERN Linux 7.3.1 Certified

Jan Iven, IT/ADC

CERN Linux 7.3.1 (a customized version of Red Hat 7.3) has been certified. This version replaces the previous CERN Linux 7.2.1 as the default supported version of Linux at CERN.

CERN Linux 7.3.1 is available for network installs and on CDs (which will be distributed through the IT bookshop (<http://consult.cern.ch/books/>). A matching CD

with ASIS packages will be made available.

Please see the URL <http://cern.ch/linux/redhat73/> for installation instructions, more documentation and an overview of the certification process itself.

We want to thank all participants, especially the members of the (newly established) Linux certification coordination body.



3.2 Solaris 8 Certified

Ignacio Reguero, IT/PS

The Solaris 8 certification process finished on the 13th September 2002, thus making Solaris 8 the recommended Solaris version at CERN.

The operating system level is Solaris 8 HW 02/02 with the latest Sun recommended patches as available in

</afs/cern.ch/project/sun/solaris/patches2.8>.

The recommended post-installation is described by the CERN SUE feature as available in

/afs/cern.ch/project/sue/pro/sun4x_58/dist/feature/cern.

The default compilers are the Sun SPARCcompilers 6.1. The SPARCcompilers 6.2 and Sun ONE Studio 7 are also available, in particular you may find the production version of Sun ONE Studio 7 with the latest compiler patches in

</afs/cern.ch/project/sun/solaris/opt/SUNWsp7>.

A SUE feature to automate the installation of Java 1.4.0 as distributed by SUN for Solaris 8 is available. This is Java(TM) 2 Runtime Environment, Standard Edition (build 1.4.0-b92) Java HotSpot(TM) Client VM (build 1.4.0-b92, mixed mode). This feature allows choosing between a remote installation using pointers to AFS and a local installation using Solaris packages. The feature is installed by executing `/usr/sue/etc/sue.install java`

The persons responsible for them in a cooperative effort have performed extensive tests of the main software packages used on Solaris at CERN with the system support people. Many

thanks for the help of everybody involved. The Solaris 8 Certification page (see the URL <http://wwwinfo.cern.ch/ce/UI/Docs/sol8cert.html>) provides the detailed list of tests performed, as well as the list of products and SUE features that are available.

The SUNINST0 Solaris network installation server has been upgraded to new hardware and software (Solaris 8). The upgrade was required due to the former SUNINST0 HW being extremely obsolete. We profited to consolidate the BOOTP/DHCP server used for Solaris installations onto the same machine. This service offers fully automated and reproducible “Jumpstart” Solaris installations through the network using BOOTP/DHCP protocol for Solaris 8 and RARP for older Solaris versions. Please contact Solaris.Support@cern.ch in order to define machines for installation.

Please note that, triggered by a request from CMS, we are working to upgrade the SUNDEV public Solaris service to Solaris 8 soon.

Should anybody have any concern or question about this upgrade please share it on the mailing list forum-solaris-certification@listbox.cern.ch.

We also propose this list for discussions on product availability and versions on Solaris, with the goal of reaching wide consensus with the Solaris community before any important move.



3.3 Tru64 Unix (Dunix) Support at CERN

Alan Lovell, IT/PS

The CERN support services for Tru64 UNIX (DUNIX) are due to be phased out over the coming months. This document details the support services that will be offered during this period.

Installation Server and Reference Machines

All support services will be stopped at the end of 2002. Until then the service will be limited to the following areas:

- Maintain currently available operating system versions without modification - OS software frozen (Version 4.0f)
- Make security patches available when necessary and when they are available.
- The installation server will be available for system (re)installations, however no new systems will be added to the RIS database
- SUE environment will be only updated to fix bugs and apply security patches.

There will be a limited general third level support offered for True64 Unix until the end of 2002. We will continue to maintain AFSoft and refdux40f, with the limitations mentioned above, during the course of this year. Otherwise the True64 support will be frozen. This service will no longer be generally available as of 1st January 2003.

Workstations

Euclid (AXCAD) alpha-stations will continue to be supported as before until the replacement program has been completed during the course of the first half of 2003.

For all other alpha-stations support will cease at the end of the year. Until then first and second line support is via the DCS contact with a best efforts third line support from IT/PS/UI. Please note that there is no formal support contract with HP/Compaq for alpha workstations other than for the AXCAD systems. This means that any intervention will be charged to the owner of the workstation. If HP/Compaq need to be called then the contract staff should inform the user that this will require payment and it would then be up to the end user to organise the intervention/payment directly with the vendor. See below for further information on vendor support.

EDMS/AXCAD Tru64 Servers support for these systems will continue as before until they have been replaced with other configurations. Plans are being finalised for the replacement of dedicated servers such as those for EDMS and Euclid. In the meantime the operating system for these systems will remain frozen on Tru64 V4.0f.

Vendor Support

Given that the version of DUNIX (4.0f), which is running on the clients and servers, has ceased to be supported as a standard product by HP/Compaq on 30th September 2002 it is essential that these systems (servers and workstations) be migrated to another platform as soon as possible. No guarantees on the stability of the service can be given as of this date although one could expect it to run for a limited period with few problems. Please note however, that as of the date when the support from HP/Compaq has ceased any system errors would not, or at least not normally, be investigated by the company unless we undertook to upgrade to a later release. Later releases of the TruCluster software have not and will not be certified in the CERN environment.



4. Internet Services and Network

Everything related to Internet in a large sense: i.e., network issues and performances, all technical Web issues, electronic mail, Internet news, home access to CERN facilities, video conferencing and multi-media, network security, etc.

4.1 A New Infrastructure for the CERN Electronic Messaging Service

Michel Christaller and Emmanuel Ormancey, IT/Internet Services

Since 1998, the Electronic Messaging at CERN has been centralized with an infrastructure that has grown to a service holding 14,000 mailboxes stored on half a Terabyte of mail data, and routing over 100,000 messages per day.

The service has been successfully operated for many years and a new infrastructure is being deployed with the goals of improving both security and user's satisfaction. The new service will now run on commodity hardware and will also offer new functionalities such as a Web-based mail interface and a shared calendaring service.

After a study phase to design a new scalable infrastructure and a pilot phase to validate the basic choices, the new service will be in official production on December 1st, 2002.

The new infrastructure will provide the same services as before (routing mails through SMTP protocol, accessing mailboxes through IMAP/POP protocols) and thus will be compatible with the current mail applications like Pine, Netscape or Outlook. The migration will be done almost transparently.

Users of the new service will also be able to benefit from:

- The new Web interface to access mail,
- Same username and password across Mail, Web and NICE services,
- Secure and encrypted sessions to the mail servers,
- A calendaring service with shared access,
- Collaborative tools like delegation to secretaries,
- Online / Offline features for travelling users,
- An enhanced spam (unsolicited mail) filter.

To benefit from some of these enhanced features, the new Web interface or a new version of Outlook (*Outlook 2002* also named *Outlook XP*, which is part of the new Office 2002 suite, available to NICE 2000 users) will be required.

The migration phase will take probably more than one year to complete, but it is already open to everyone. Interested people can just have a look at the new Mail Service web site (<http://cern.ch/mmmservices>), where a form can be filled in to ask for migration. Anyway a direct invitation to migrate will progressively be sent to every CERN Mail account.



4.2 Receiving Unwanted Email?

Denise Heagerty, CERN Computer Security Officer

Abstract

This article gives advice on recognising SPAM (unsolicited email) and viruses so you can delete them quickly.

Unsolicited email, commonly known as SPAM, is a growing problem across the Internet at large. You are recommended to *delete it*. If the mail looks suspicious, *do not even read it* (from Outlook you can delete email using the right button of your mouse to select the message and then selecting delete). Viruses are often hidden inside attachments, so *do not open unexpected attachments* (if you are unsure check with the sender first). SPAM and virus emails can be disguised to trick you into reading the email and/or performing an action. Here are examples of some techniques to help you recognise them:

- *faked email addresses*: mails can appear to be from people you know or even from yourself. Your own email address can appear in mails which you did not send, resulting in non-delivery messages or unexpected replies.

- *enticing subjects*: the mail subject uses words to make you curious, believe the email is important, or specific to you, so that you will read it.
Opening an attachment could infect your computer with a virus.
- *asking you to forward email to people you know*: this is probably a viruses or false information - do not forward such email.
- *join a petition or support a cause*: the petition or cause rarely exists, it is more likely that your email address will be collected and used for further SPAM mails (as sender or receiver). References to recent or topical events are common techniques to make the false information look more realistic.
- *click on a web site*: if you click you could be downloading a virus. Click cancel (instead of ok) or close unexpected dialogue boxes when using the web. This can also be a technique to validate your email address and increase your chances of receiving more unwanted emails.
- *money scams*: one of many examples of false informa-

tion — do not believe them just delete them.

- *response to your request*: the mail may pretend that you made a request so that you will take it more seriously.
- *remove from a list*: asking you take action to remove yourself from a list you did not join can be used to validate your address for (ab)use in the same way as petitions and causes.

CERN's email gateways filter more than 45,000 SPAM emails every day, but still such email can reach you as dis-

tinguishing SPAM from intended email is not easy. Further advice and details of CERN's anti-SPAM techniques are at URL:

<http://cern.ch/MailServices/docs/problems/spam/spam.html>

CERN's computer security recommendations also provide useful guidance:

<http://cern.ch/security/recommendations>



4.3 Batch Nodes on Non-Routable Network

Ulrich Fuchs, IT/FIO

Preparing for the LHC computing era and its demands, we are now addressing a problem that has already been around for some time: the shortage of IP numbers.

Already today a free IP number is a precious thing at CERN and planning the massive machine installations for LHC computing, a practical solution had to be found. Consequently new Lxbatch cluster nodes will be installed on a private (non-routable) network, starting from the 2003 acquisitions. Additionally, machine room renovation work will necessitate moving a certain fraction of the current batch machines into the vault, where they will be re-installed and moved to this private network over the next weeks. Please note that this applies only to the batch worker nodes, other nodes such as the gateway nodes for the Grid will naturally be on routed networks.

What is a “private (non-routable) network”?

The worldwide IP network community agreed on certain IP number ranges that are not supposed to be routed outside LANs. This means that a host with one of these numbers can only be seen by its immediate neighbours in the same network segment. All network traffic caused by this host is removed from the network by the first router on its path and so is any traffic aiming at this host from any other machine worldwide.

What are the implications at CERN?

The CERN backbone routers were re-configured to handle private IP traffic within CERN but will not route packets to the “outside world”. A host running on private IP is therefore visible to all hosts on the CERN network and can be used as any other machine (from inside CERN) but *cannot be contacted from outside* and also *cannot contact any machine outside CERN*.

If your batch job is running on a machine on private IP it has full access to all resources at CERN but cannot contact the outside world by any means (no ftp, scp, ssh, AFS, NFS, ...). The majority of batch jobs running today only access

resources either on the node itself or at CERN (Castor, the CERN AFS cell). For these jobs it does not matter whether they run on a private network or routable network. Only jobs accessing resources outside CERN need to run on machines on the routable network.

What can you do?

To enable you to test your programs in a private network environment, we have set up 2 machines (one Linux6 and one Linux7) on this network already and provide a special LSF queue called “wan_test” to access them. Please feel free to use this queue and test your software already today.

Wouldn't using IPv6 provide all the routed addresses we need?

There are some IPv6 pilot projects at CERN but we're still years away from a production-ready service.

What about the impact on Grid computing?

Worker (CPU) nodes are foreseen to be on non-routable and gateways on routable networks.

Will private non-routable IPs also be available to “users”?

Private IP non-routable subnets can be given to certain projects, IT/CS group will happily discuss your needs.

How can one select specific nodes if one needs WAN access?

Hosts on the routed network have a resource “wan” set. When there are production nodes on non-routable network available to the shared public queues jobs will be dispatched equally to WAN and non-WAN nodes. So if your job needs external access you have to submit it with the -R wan option, otherwise it will go to any host.

How many nodes will be on either network?

All new nodes installed from now on will be on the non-routable network. Also the nodes that will move to the vault this year will change to the non-routable network (i.e., by March 2003 more than 60% of the total batch capacity).



5. Scientific Applications and Software Engineering

This includes descriptions of: scientific applications and libraries from CERNLIB and Anaphe, mainly used by the Physicists; engineering applications like the ones on PaRC or those associated to online controls; software development environment.

5.1 Anaphe – OO Software for HEP in C++ and Python

Andreas Pfeiffer, IT/API

Abstract

The Anaphe project provides an Object Oriented software environment for HEP experiments. Based on public domain, HEP specific and commercial software a suite of C++ class libraries implementing the components of the AIDA Abstract Interfaces have been developed. A thin layer allows to access this functionality from the Python interpreted language providing a powerful tool for interactive data analysis.

Introduction

The Anaphe project provides an Object Oriented software environment for HEP experiments. Public domain libraries cover the basic functionalities; on top of these libraries a set of HEP-specific C++ class libraries for histogram management, fitting and ntuple-like data analysis has been developed. In addition to the license free public domain version, the functionality can be extended by using commercial libraries such as NAG-C.

In close collaboration with the experiments a new modular data analysis tool/toolkit is being developed. Based on AIDA compliant Abstract Interfaces, a set of libraries has been developed to implement these interfaces, making use of already existing software (e.g., from public domain and commercial libraries) wherever possible. In order to comply with the user requirements for a command-line driven tool, a “mapping” of the Interfaces to a scripting language (Python) is provided.

The loose coupling provided by the consequent use of Abstract Interfaces for each component in combination with the use of shared libraries for their implementation allows the integration into modern scripting languages. This integration is simplified even further by using a specialised toolkit (SWIG) to create “shadow classes” for the Python language, which map to the definitions of the Abstract Interfaces almost at a one-to-one level.

AIDA — Abstract Interfaces for Data Analysis

The goals of the AIDA project are to define abstract interfaces for common physics analysis objects, such as histograms, ntuples, fitters, I/O, etc. The adoption of these interfaces should make it easier for physicists to use different tools without having to learn new interfaces. Additional benefits will be the interoperability of AIDA compliant applications, for example by exchanging analysis objects via XML between applications. The use of the AIDA interfaces in the

user’s code also eliminates the need for changes in the user’s code when changing from one implementation of the AIDA interfaces to another. Within AIDA bindings to AIDA for both C++ and Java have been developed.

Finally by bringing together the authors of a number of different applications we expect to gain deeper insight into the strengths and weaknesses of different applications, and to improve collaboration between different authors.

Several implementations of the AIDA interfaces exist today and are used by a growing community of active users.

Anaphe and Lizard

Anaphe – the toolkit

The package consists of a set of *foundation libraries* and a set of libraries which implement the *AIDA interfaces* using functionality from the foundation libraries. In turn, some of the foundation libraries rely on functionality provided in a lower layer of the software.

This lower layer includes open source software (e.g., Grace, Qt, SWIG, GSL, Python, expat), commercial software libraries (e.g., OpenGL, OpenInventor, NAG-C, Objectivity/DB), older CERN software (e.g., HBOOK, Minuit) and software provided by other projects (e.g., CLHEP).

The next layer is implementing the AIDA interfaces in separate components to allow decoupled development and maintenance. Some of the functionality actually needs alternative implementations: examples are the *fitter* component where the underlying minimizer engine can be selected to be Minuit or from NAG and the *store* component where AIDA objects can be stored/retrieved from HBook files, (compressed) files with XML formatted objects, and Objectivity/DB databases.

The component based approach we have chosen allows for an easy extension of any given component without the need to modify the other parts (components) of the system.

Lizard – the tool

Since July 2001, we have also provided an interactive analysis and visualisation tool called *Lizard*, which provides a flexible framework which can be fully configured at run-time. This interactive tool is based on Python, a powerful and intuitive open-source OO interpreted language which is very widely used in a variety of domains.

Lizard is actually a very “thin” layer on top of the Anaphe libraries: essentially it is a mapping of the AIDA interfaces to Python classes; the Anaphe libraries implementing the AIDA interfaces are loaded at start-up time dynamically. Thus we profit from the flexibility of selecting the desired specific implementation (e.g., Minuit minimisation engine for the fitter component) at start-up time without the need for re-linking the application.

Adding experiment-specific libraries to Python is surprisingly quick and easy. Lizard may be easily extended with user-defined functionality or any of the growing list of Python modules. Of course, “kumac-style” Python scripts may be written to automate repetitive tasks or build more complex analyses inside the Lizard environment. You can also write your own Python “shortcuts” to compress more complex Python procedures into single commands.

Arbitrary C++ code can be compiled, linked and run on-the-fly. Access to the histograms, ntuples and “vectors” (arbitrary data points) inside Lizard from this code is transparent.

Usage and download

The actual release of Anaphe (version 5, which implements the AIDA interfaces version 3.0) is available for RedHat61, RedHat72 (and 73) with compilers gcc-2.95.2 as well as for RedHat73 with gcc-3.2 and Solaris-8 with CC 5.4. Users have reported that it also works well on SuSE 7.2/gcc-2.95.3 (use the redhat61/gcc-2.95.2 version) and SuSE 8.0/gcc-2.95.3 (use the redhat72/gcc-2.95.2 version).

There are basically four scripts provided to use AIDA/Anaphe/Lizard in your environment:

- `lizard`: to start the interactive tool Lizard;
- `aida-config`: a script to return the flags needed in compilation and linking programs with an implementation of AIDA, similar like the `cernlib` command (`aida-config -inc -lib`);
- `setupAnaphe` (and `setupAnaphe.csh`): file to be *source'd* to prepare the run-time environment (setting environment variables like `PATH` and `LD_LIBRARY_PATH`) before running an executable program which uses AIDA through Anaphe.

From a machine which has access to the CERN AFS cell (e.g., on lxplus or lxplus7), simply execute the corresponding script from the following directory:

```
/afs/cern.ch/sw/lhcxx/share/LHCXX/latest/scripts/
(you can replace “latest” by a specific version number
if you wish). For a more convenient use, you might want to
set a link from your $HOME/bin/ directory to these scripts:
s=/afs/cern.ch/sw/lhcxx/share/LHCXX/latest/scripts
ln -s $s/* ~/bin/
```

where the `s` variable was used for convenience. You then can simply call the scripts by their name (no directory prefix needed).

If you want to install Anaphe version 5 locally on your machine, please download the script from our download

page and follow the instructions provided there. The download script will create four scripts in your `$HOME/bin/` directory allowing you to use `lizard`, `aida-config` and `setupAnaphe` without explicitly specifying their path.

A screenshot of a Lizard sessions is shown in Figure 5.1 on the next page. The actual implementation of the Plotter component is based on Grace, a very powerful plotting tool which allows direct manipulation of all the elements in the canvas interactively (simply double-click on the part you want to modify or use one of the pull-down menus).

The Anaphe project is a suite of Object Oriented software covering a wide range of functionalities, from generic basic classes to HEP-specific C++ class libraries for histogram management, fitting and ntuple-like data analysis implementing the AIDA standard interfaces.

The loose coupling provided by the consequent use of Abstract Interfaces for each component in combination with the use of shared libraries for their implementation allows the easy integration into modern scripting languages like Python; a very “thin” layer on top of the AIDA interfaces provides the base for the Lizard analysis tool using the C++ class libraries in Anaphe.

Re-using existing tools and libraries from public domain or commercial providers allows to further reduce the maintenance effort while leveraging some of the resources available outside HEP.

With the start of activities in the LHC Computing Grid (LCG) project, further development work for Anaphe has been suspended (support and maintenance will continue), allowing for an alignment of the activities of the LHC experiments and EP and IT divisions at CERN towards the development of common solutions in the applications area. The architectural blueprint report defines the architecture of such common software to be component based with Abstract Interfaces as the only means of communicating between different components. It further recommends the use of AIDA Interfaces, Python and Qt; so it seems likely that future common analysis components will follow a similar style as presently in Anaphe. It is therefore expected that the users of Anaphe developing their applications based on the AIDA interfaces would also interoperate with components developed in the context of LCG in a fully transparent way.

References

- <http://aida.freehep.org/>
The AIDA web page also contains links to all products implementing the AIDA interfaces.
- <http://cern.ch/Anaphe>
The Anaphe web page also contains links to other products of the Anaphe suite.
- <http://cern.ch/lcg>
The web site of the LHC Computing Grid project.
- <http://lcg.web.cern.ch/LCG/SC2/RTAG8/finalreport.doc>
The architectural blueprint report of the LCG application area.




Figure 5.1: Screenshot of a Lizard session



5.2 2002 CERNLIB Release

Ian McLaren, IT/API

The CERN Program Library 2002 version became the “pro” version on 16 September 2002. Stdhep version 5.01 was added on 22 October 2002 and a minor update was made in graflib and PAW on 13 November. As agreed in Focus meetings last year this will be the final “pro” release of Cernlib although there will be support till the end of 2003. Details of these plans and the meetings that produced them are in a 2001 CNL article (<http://cern.ch/ref/CERN/CNL/2001/002/cernlib-plan/>). Originally intended for the spring, the release date was decided by the introduction of the Linux 2.4 kernel used by Redhat 7.x and most other Linux suppliers, which required modifications to the 2001 sources. However, the change most likely to affect existing user applications is the separation of the Castor (alias shift/RFIO) interface routines from packlib and kernlib.

Major changes

- The shift(rfio) package has been separated from packlib and requires libshift.so or its static equivalent libshift.a, mainly to avoid Cernlib getting behind developments in Castor. At CERN, the shift libraries are usually installed in /usr/local/lib and it is sufficient to add -lshift after the Cernlib libraries in the link command. Failing to include this library will typically produces errors like

```
ld:
Error: Undefined:
serrno
rfio_errno
```

For the convenience of external sites, versions of packlib and kernlib with the default system i/o routines are provided in the compressed tar files, and there is also

a noshift version of the `cernlib` command. The default configuration files are also for the non-shift version as this seems more appropriate and simpler for users building their own versions.

- Geant321 0.14 has some fixes and larger arrays.
- Paw, Hbook, Hplot and Higz: many minor fixes.
- There are new versions of the Monte Carlo libraries stdhep 5.01, isajet758, pythia6205, ariadne412, leptog651 and herwig64.

Release Notes

Access to all the Cernlib material is via the URL <http://cern.ch/cernlib>.

Full details of the changes with respect to the 2001 release are contained in the release notes (<http://wwwinfo.cern.ch/asd/cernlib/news/cernlib.news2002.html>), including system and compiler levels and a list of all the modified routines.



5.3 EXtreme Programming (XP): Pragmatism & Quality

Eric Poinsignon, IT/PS

Abstract

A short introduction to XP, one of the best-known methodologies but the most controversial.

Introduction and Characteristics

In the set of the agile methods, XP is one of the best-known methodologies but also the most controversial.

XP is very pragmatic: some aspects are naturally intuitive (like short iteration, simple conception, fast feedback and communication) but others are not approved by the developers community (like pair programming). In industry, one can see XP is not applied as a whole and it needs experienced engineers (who already followed classical methods).

There is also a high dependency on the context. For example, the Planning Game (game where developers and client plan dates of versions) and the Adaptive Change Management can only be applied in a technical assistance, but not in a fixed price contract.

The 4 values of XP

1. *Communication*: it must appear everywhere and at all times between developers and client (specification and test), between developers and manager (estimation and test), between developers (pair programming).
2. *Simplicity*: XP encourages choice of the simplest solution first (maybe later increase the complexity). This

idea is to be applied to the conception, to the planning and even to the team.

3. *Feedback*: The client must be able to detect any gap in the planning or in the features.
4. *Courage*: The client must be courageous enough to prioritise his needs and admit his unclear needs. The developer must be courageous enough to modify his architecture, even late in the process.

Papers and references

I do not want to go into detail on XP as many interesting articles are accessible on the web. So please have a look at the *Programming and Agile Methodologies* presentation of M. Michele Marchesi, who visited CERN on July 10th 2002. (See <http://cern.ch/computing-colloquia/past.htm> or <http://sdt.cern.ch/RUP/Doc.html>, the section "EXtreme Programming").

The Software development Tools Service (SDT) at CERN is providing the Rational Unify Process (RUP) where you can find some articles about interaction between RUP and XP (<http://sdt.cern.ch/RUP/>). You can find interesting "other sites" on the SDT page, at the URL <http://sdt.cern.ch/RUP/Doc.html>.

For questions you can contact the SDT service at sdt.support@cern.ch



6. The Learning Zone

All kinds of “educational” computing issues, not necessarily directly CERN-related, such as computing news from the Press and tutorials on various subjects.

6.1 Questions and Answers from the Computing Helpdesk

Collected by the User Assistance Team in IT/US

Nicole Crémel (Editor), IT/User Services

Abstract

This is a collection of Questions & Answers that have been treated by the Computing Helpdesk managed by the User Services group.

N.B. The number in parentheses refers to its relative numbering in the “Question-Answer” database, at the URL: <http://cern.ch/consult/qa/NNN>, where NNN is the problem identifier (number).

Windows Support

Question [3083] – NICE W2K boot floppy creation doesn't work

I have followed the instructions concerning the creation of a CERN W2K Diane boot floppy from <http://winservices>, section NICE 2000 installation.

The file download to disk works ok, but when I try executing the file, a DOS window pops up a couple of seconds and disappears again. And that's it...

Answer

This is a known but not yet fully understood bug which has been seen on several computers. The easiest way out is to either reboot or simply use another PC in order to correctly produce such an installation floppy.

A few cases have been seen where the problem seemed to be related to the PATH environment variable on the machine in question. In this case, the user should execute:

Control panel

- Add/Remove Programs
- Add New Programs

select the category 'Utilities' and look for

CERN - Correct PATH Environment Variable

and click "Add".

If this facility is not visible in Add/Remove programs, the use must logout/login to get it presented.

Question [3096] – Questions for laptop NOT connected to NICE

I have a laptop which is registered in the CERN database but does not join NICE. I have a few questions of principle:

1. Is it possible to see other CERN NICE/PC machines as part of the network neighbourhood?
2. Is it possible to see files in NICE home servers?
3. Is it possible to print on CERN printers?
4. Should all above not be possible: Can I join the NICE domain but not have the automatic NICE software updates?

Answer

Here are the reply to your questions:

1. Yes, but the computer you want to access should have access rights for you. Your friends/colleagues will probably give you such rights.
2. CERN will not allow you access to its servers (e.g. cernhome01). Since you do not join the CERN domain, you may not have all the security packages installed, nor have all the anti-virus updates.
3. All information about the CERN printer wizard, including how to download it can be found at <http://cern.ch/PrintPackage>
4. This is not possible. If you join the CERN domain you automatically get all the security, virus killer, Office and other future updates.

Conclusion:

You may want to consider joining the NICE domain as this will be the answer to all above questions. However, by joining the CERN domain you automatically get all the security, virus killer, Office and other future updates.

Question [3100] – Multiple clipboard items in Office2000

I sometimes need to manoeuvre between two or even more items that I previously copied into my PC's clipboard. Is there any way to find entries older than the most recent one somewhere, or is there any other way to handle multiple copy/paste operations?

Answer

This is only possible on Office 2000 (or higher), however, unfortunately not for any other W2000 application like NotePad, Exceed, Remedy or whatever else.

Open an Office2000 application. You will see a taskbar displayed, containing fields like "Paste all", "Items", etc., provided you have ticked View ? Toolbars ? Clipboard.

The *Items* collection contains the last few (up to 12) clipboard entries. Moving the mouse over any of them, displays its first 50 characters. Click to insert. The rest of these few symbols is self-explanatory.

Note that entering more than 12 clipboard items will get rid of the oldest one(s).

Question [3103] – Reduce size of bitmap files

Is there a tool available on NICE 2000 that I could use to reduce the size of my bitmap files, without losing quality?

Answer

Bitmap files (extension ".bmp") are quite space consuming. I suggest you convert them into ".gif" or ".jpg" files with *Paint*.

The resulting files will be smaller, which means they will use less memory and disk, and load quicker.

- Call up *Paint*:
Start - programs - accessories - *Paint*
- Open the .bmp file
- Select the file
- Save As: choose .gif or .jpg file type explicitly

Your file is now considerably smaller in size and can be read/edited by other appropriate software, for instance *Microsoft Photo Editor*.

Question [3107] – Edit screen shots on NICE2000

Can I crop/edit a screen shot before I definitively paste the bits I need into an Office2000 document?

Answer

Screen shots are in bitmap (.bmp) file format. One of the tools available on NICE2000 to edit them, is *Microsoft Photo Editor*:

- Make your screen-shot the usual way (PrintScreen for a full shot, or ALT+PrintScreen for the window only where the cursor sits).
- Start - Programs - Microsoft Office Tools - Microsoft Photo Editor.
- Select the Edit menu, click Paste as New Image.
- If you only want to select a portion of your screen shot (the portion you want to copy to your Office document), click Select on the Standard toolbar (the symbol that looks like a dashed square), then drag the cursor over the area you want to select.
- Select the Edit menu, click Copy.
- In the Office document you want to copy the selection into, do an Edit/Paste.

Question [3112] – Disk cleanup feature on W2000

On my NICE2000 PC, I see a feature called "disk cleanup":

- Start - programs - accessories
- system tools - disk cleanup

What exactly does it do, and am I safe to use it just as it is?

Answer

Disk Cleanup safely gets rid of all temporary, unnecessary, doubled, etc., locally stored files that you probably never need anymore. It is part of W2000 and acts a bit like earlier, similar 3rd party software (for instance "*Norton Clean-Sweep*" and the like).

It can only be used on local NTFS drives. These are analysed one at a time, after which a menu of cleanup options is presented to you. It lets you decide what to delete, like for instance temporary Internet files, off-line files etc.

A nice feature it also offers is the possibility of compressing old files.

You have to specify the age of the files as of when compressing should be done. This does not delete them, but just reduces them considerably in size. Make sure compressed files are marked as being displayed in an alternate colour afterwards, so tick this option in:

My computer - tools - folder options - view

Note that such disk cleaning operation takes quite some time and slows down the computer performance.

Question [3113] – Usage of share folders

I have created some local W2000 shares on my hard disk, that I enable for quite a few user accounts of my working environment.

How can I find out, e.g., in case of a deadlock situation, which colleague uses which file?

Answer

- Right-click "My Computer"
- Choose "Manage"
- Open the "Shared Folders" tree
- Click on "Open Files"

The right/hand menu shows the user account and the file(s) that (s)he has open on your share(s).

This information will also be useful in case you want to reboot or run some local backup.

Question [3116] – Hide shared local folder names

I share a few of my local W2000 folders with some colleagues. However, I noticed that all other CERN users who can see my computer over the network, can also see these folder names (although they cannot access them, of course).

This can be undesirable, as I don't even want others to know of the existence of such shared folders: even their names might sometimes be confidential. Can I hide them somehow?

Answer

There are two possible solutions:

1. Create the shared folder with a name that ends with the "\$" (dollar) sign. It will then be invisible, but to *all* network users, including the ones that have access granted. Therefore, these users will have to address the folder by its exact full path name (e.g., \\pcjoepublic\charm-quarks\$).
Note: renaming an already existing share loses all its access controls which you will have to recreate afterwards! A Windows prompt will remind you of this.
2. Just make the folder you want to hide a sub-folder within another one which carries some neutral name. Unauthorised users will then only see this 1st level name, but won't be able to access any further down its tree.

Question [3123] – NICE trustee not changing ACLs

The NICE trustee manager does not change any access rights of a folder, although I am empowered to do so with my account that is listed as 'access control enabled'.

Answer

Someone (probably with administrator rights) has deleted the access control entries for CERN\Domain Admins and/or BUILTIN\Administrators.

Note: these ACLs should never be touched as they are vital for the trustee mechanism.

The Trustee Manager is running as a domain administrator. If the ACL of a folder does not contain the Administrators group, the Trustee Manager has no privileges to change anything. This happens when you manually remove such rights. You should also be aware that even the backup cannot work under these conditions!

Only way out is to forward this case to a Windows system responsible (via Nice.Support@cern.ch).

Question [3137] – Replace existing file - option NO to ALL

I am working with lots of files which I occasionally transfer into other folders. However, I do not want to replace already existing files of the same names, with new ones.

Windows offers me (apart from cancel) three choices when it says:

The folder already contains a file named 'file.ext'

Would you like to replace the existing file with this one?

YES - YES to ALL - NO

Answer

The function No to All is well hidden. To activate it, click No with the SHIFT key pressed at the same time.

Question [3099] – Word: Multiple page layout

I have a Word2000 document that contains a few hundred pages. How can I work in multi-page mode, i.e. how can I view and edit some pages occasionally, without having to look at them one by one?

Working my way down with the scroll bar is rather clumsy...

Answer

Word offers a nice feature, letting you display and edit multiple pages within one layout. This is specially useful for pages with lots of graphics, but of course also for simple text, when one doesn't know or care which is the exact page number to edit.

Open the Word document and:

- view / print layout
- view / zoom
- tick "many pages"
 - go with the cursor into the little PC monitor symbol, left-click and drag. All possible multiple page combinations will be offered, between "1 times 2 pages" and "4 times 9 pages". Choose the one you want, then let loose the mouse button.
- After clicking OK, Word will now display a multi-page page.

Unix Support

Question [3032] – Request for large amount of space on Unix (AFS)

We need quite a lot of space on Unix to run our software (ATLAS collaboration). Is it possible to get some kind of "maxidisk"? Can such a maxidisk be used (and shared) by several users?

Answer

Please note first that there are three major sources of disk space on Unix central servers:

1. Local temporary space (e.g., /tmp) on your machine. This is probably the best for most types of access.
2. Global space (at CERN it is /afs/cern.ch using AFS) which has a caching mechanism and a central server. This is ideal for moderately small amounts of data and source code. It is backed up nightly, and has the convenience and risk and extra cost associated with these features.
3. Data mass storage (i.e., SHIFT/Castor at CERN).

AFS space is managed per account (with a maximum 500 MB) or project directory which may be personal or experiment wide.

If you think the solution for you is AFS space, it is possible to create a 500 MB maxidisk, e.g.,

/afs/cern.ch/atlas/maxidisk/d141. Such a request must be addressed to your experiment mailing list (e.g. atlas-support@cern.ch), or to a computer administrator for ATLAS (group ZP, look at: <http://consult.cern.ch/xwho/people/form/admins>).

Access to this space is governed by the usual AFS rules, and it is possible to configure it (e.g. give read or write access to several users) with the usual command `fs setacl` (`fs listacl` will display current access and ACLs setting).

Question [3188] – Scratch space for batch (LSF)

What directories should I use for batch jobs with LSF (like a scratch disk) for temporary storage of files?

Answer

Each batch node on which we run LSF jobs has a big /pool partition, to which LSF execute the `cd` command before your job starts. We are trying to make this partition as big as possible. Therefore you should use this space to store your big amounts of data.

Note that your AFS home directory is available during your job execution, for read and write, because LSF provides you with an AFS token. In addition, some experiments have AFS scratch space, which is normally of the order of GB, to store data. Check this with your AFS administrator of your group.

If you want to keep your data in /pool after the job has finished, you should use our CASTOR system to store it on tapes, because this is the system that can handle big amounts of data, much better than AFS or even NFS. We do not have scratch space on NFS. You have to do the copy within your batch job, e.g. as the last command, because you are not allowed to log in to the batch nodes after the job has finished, and, in case the next job that comes along after yours, wants to have a big /pool file system as well. On batch machines you should not use /tmp for big data storage.

To learn more about CASTOR, please refer to the WEB pages:

<http://cern.ch/it-div-ds/HSM/CASTOR/Welcome.html>

Question [3179] – Remote X Applications between 2 Linux machines (RedHat 7.2)

Having recently installed CERN RH 7.2.1 on my Linux PC it seems that the X11 forwarding does not work: I cannot start X applications when opening a remote connection (with ssh) on another Linux PC at CERN.

Answer

Most probably this is because of the firewall which is active by default on Linux PCs at CERN.

We recommend to use `mxconns` in this case, which will provides some safeguards against active malicious X11 requests. It is explained at:

<http://cern.ch/linux/redhat72/documentation/faq/index.shtml#mxconns>

Note: system administrators can change the configuration of the firewall on their PC with the command

setup / Firewall configuration.

Mail Support

Question [3076] – How do I add an additional GEM?

How do I create, in addition to the generic e-mail I already have been given, another GEM?

Answer

This can only be done via a special request. So, please send an e-mail to User.Registration@cern.ch, specifying your name, physical e-mail address (PEM) and the additional GEM you would like to get. Obviously, you cannot get a GEM which already exists for another person.

Be prepared to wait at least one day until the new GEM works.

Please note: we strongly recommend to *never change* an existing GEM!

Question [3184] – Invalid URL when reading a mail with Outlook

I received a mail pointing to the URLs: <http://documents/Convert> and <http://consult/qa/0377>

However, when reading this mail on my NICE PC with Outlook, I cannot open those two URLs. It seems that Outlook is automatically adding a ' .' (dot) at the end of the URL, making the link invalid.

Answer

This is – apparently – a bug with Outlook 2000 on Windows 2000. We can only propose some “workaround” but only at the time of creation of the message that contains such a URL, where you should:

- always include "cern.ch" (or the domain) in the URL (e.g., <http://documents.cern.ch/Convert>).
- add a ' .' at the end of the URL.
- instead of using plain text or rich text format, format your message as HTML.

N.B. This trouble is not present on Outlook 2002 (known as Outlook XP).

Question [3070] – Where is my mailbox? (mail server)

We sometimes are informed that one out of the many mail servers has problems, or will be down for some time.

How can I find out which of the CERN mail servers I am served by?

Answer

You can go to the web site <http://cern.ch/mmm> and open the link Where is my mailbox?.

Then just type your mail login, and the information will be returned (this only works inside CERN).

Another solution is to use the command

nslookup xxx.mailbox,

where xxx is your mail loginid.

N.B. On a Windows PC, you can access this command by opening a cmd.exe window with:

Start - Run... - cmd

(Command is located in C:\WINNT\System32). and then execute nslookup xxx.mailbox

Web Support

Question [3134] – Start a program from Internet Explorer

My colleague is able to launch a programme directly from the address field of Internet Explorer on Windows 2000, by just typing its name. When I try the same, IE directs me to an MSN search engine without result.

Answer

Internet Explorer can launch any executable programme, provided it finds its shortcut on the desktop. You need first to create a desktop shortcut of the program you would like to launch directly from IE. Its exact name must be used afterwards, so keep it as short and meaningful as possible.

For programs that you often use this method is quicker than having to shrink some Windows and double-clicking the desktop item, let alone looking for the programme in the 'Start - programs' list.

Question [0351] – Cache in Internet Explorer

How to clear the cache in Internet Explorer?

Is it possible to define the size of the cache (in order to limit the amount of data cached)?

Answer

In both cases select in the menu Tools:

Internet Options...

In the panel General, for Temporary Internet files:

- to clear the cache select Delete Files... and also Clear History (into History).
 - to define the size of the cache select Settings...
-

Question [0020] – Electronic encyclopedia on the Web

Is there a facility on the Web to get all kind of technical information (e.g. an electronic encyclopedia)?

Answer

You can get access the electronic version of the *Encyclopedia Britannica* from the CERN Home page following the links: "Library & Archives" (for "Events and Publications") and then "Britannica", or directly at URL: <http://www.eb.co.uk:180/>

It gives access to a list of "articles", an "index" or a "dictionary" (by selecting the appropriate square box).



7. User Documentation

Everything related to “computer documentation”: updates to computer documentation distributed by IT, offerings of the Computing Bookshop, etc.

7.1 News from the User Services Bookshop

Jutta Megies and Roger Woolnough, IT/User Services

Abstract

Since it is sometime since the last news, there are many new titles and editions available. You will find them listed below.

The Particle Odyssey - A Journey to the Heart of the Matter

In conjunction with ETT we have a new particle physics book co-authored by Christine Sutton and Frank Close. It is now available from the shop or reception in building 33: 240 pages lavishly illustrated in colour and full of information about CERN and various experiments. The book explains fundamental discoveries which have led to our current understanding of the origins and nature of the material universe. A century of particle physics, the laboratories and personalities are represented as well as many photos of particle ‘events’ (price is 60ChFrs).

New books and new editions

Other new books include an addition to the Python range from O’Reilly, “*The Python Cookbook*”. There is a new C++ book, “*Modern C++ Design*” and we have added both the CD and examples book in the range for “*Numerical Recipes*

in C++”. There are also new books on SSH and Network security.

Here is a list of the books that we have in new editions:

- *XML in a Nutshell*
- *Dynamic HTML*
- *HTML & XHTML the Definitive Guide*
- *Linux Kernel Programming*
- *Developing software with UML*
- *Oracle PL/SQL Programming*
- *Perl 5 pocket reference*
- *Essential System Administration*

Linux CDs

In the coming weeks we will be able to sell the latest version of Linux certified at CERN, version 7.3. As it was not possible to copy the CDs for the last version this time it will be a 3 CD set and should be in the shop by start December. A computing news will be sent when the CDs are available.

If you want to contact us

Any suggestions for new books are welcome, please mail bookshop@cern.ch.



7.2 User Services Book Catalogue

Jutta Megies and Roger Woolnough, IT/User Services

Situated in Building 513 1-022 (tel. 74050) is the Computing Bookshop provided by the *IT/User Services* group, where CERN users can find computer books and CDs at discount prices. The service is open weekdays from 8.30 to 12.30 or contactable by e-mail to Bookshop@cern.ch. Books are purchased from some 15 publishing houses and the catalogue offers a selection of documentation aimed at the range

of computing utilities available at CERN. The service welcomes suggestions from the user community for new acquisitions. Purchasing may be done internally via EDH or TID or alternatively cash payments via the CERN Bank.

The list, with all the relevant information, is compiled regularly at URL:

<http://cern.ch/consult/books>.

