


CERN COMPUTER NEWSLETTER

CREMEL
Nicole Laurence

RES-S/IT-US



CERN - IT Division
CH-1211 Geneve 23
SWITZERLAND

APRIL – JUNE 1999

No. 235

Contents

If you need Help	i
1 Announcements	1
1.1 New Help Desk Service in Building 40	1
1.2 Switching Off The Dial In Modem Service	1
1.3 End of the CSF/NQS Service	1
2 Physics Computing	2
2.1 Invitation to Perform Y2K Testing	2
2.2 Using a Single logonid for Multiple Groups or Experiments	2
2.3 New STK 9840 Tape Units Enter Service	4
2.4 Connecting to CERN from Home: ACB Service over ISDN	5
3 Desktop Computing	6
3.1 <i>Springer</i> UNIX Print Server Removal	6
3.2 New Print Package on NICE	6
3.3 Windows NT <i>Service Pack 4</i> Installation on NICE NT Systems	7
3.4 Y2K Compliance of UNIX Systems	9
3.5 Operating System Certification - Present List	9
3.6 Operating Systems Availability on ASIS	10
4 Internet Services	11
4.1 Configuring an Internet Access Provider	11
4.2 Computer Networking is Getting Faster and Better Organised	11
4.3 How (Not) to Block the Computer Networking and...	12
4.4 CERN's Connections to European and Global Internet Improved	13
5 Scientific Applications and Software Engineering	14
5.1 Garfield Version 6.29	14
5.2 Software Development Tools Service	15
6 The Learning Zone	16
6.1 A Tutorial on Character Code Issues	16
6.2 Questions and Answers from the Computing Help Desk	22
7 User Documentation	27
7.1 News from the User Support Bookshop	27
7.2 The UCO Reference Cards - Current List	28
7.3 User Support Book Catalogue	28

Information for Authors and Readers

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

Past experience shows that it is a good policy to have four CNL editions per year (one around Easter, one before and one after the summer, and the last one just before Christmas), unless the number of contributions is such as to make it necessary to have a supplementary issue, or to cancel one. Accordingly,

the deadline for contributions to the next issue of the CNL is:

Issue number	Deadline
236	Monday 6th September 1999.

Contributions to the CNL in HTML will be preferred, although plain text is also acceptable. Articles, news items and letters intended for publication in the next issue should be *sent directly to the editor* (Cnl.Editor@cern.ch) with no particular subject line required.

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.

This CNL issue is available on WWW at the URL http://consult.cern.ch/cnls/235

This document was produced from the HTML master, which was first transformed to XML with Tidy, and then translated into $\text{\LaTeX}$ with an XSL style sheet. To optimize the printed layout some editing might have been necessary, so that this printable version may differ a little from the original HTML source file. In any case the Web version (HTML) must be taken as the reference. A (compressed) PostScript file, containing the complete printable version of this CNL can be obtained by anonymous ftp to [asisftp.cern.ch](ftp://asisftp.cern.ch) as follows (commands to be typed by the user are underlined):

```
ftp asisftp.cern.ch
Name (asisftp:username): anonymous
331 Guest login ok, send your complete e-mail address as password.
Password: yourusername@yournode.domain.country
ftp> cd cnl/235
ftp> binary
ftp> get cnl235.ps.gz (or get cnl235.ps)
ftp> quit
```

If you do not have the gnu gunzip utility on your system you can get the uncompressed PostScript version by typing the command get cnl235.ps, without the gz suffix. To save Internet bandwidth we, however, advice you to install the gunzip utility since gzipped files are about three times smaller than their unzipped equivalents.

All **subscriptions to our distribution lists** can be made through the CNL "Help&Info" page, at URL:

<http://consult.cern.ch/cnls/help>

It is possible to subscribe:

- To the *CNL Announcement Distribution List* (<http://consult.cern.ch/cnls/help/announcement>), to receive a mail whenever a new issue of the CERN Computer NewsLetter is available on the Web.
- To the *CNL Paper Version Registration* (<http://consult.cern.ch/cnls/help/register>), to receive the paper version of the CNL at a given address.

You can also cancel your subscription in either case by filling in the appropriate form (same URL as above).

<i>Responsible editor:</i> Nicole Cr��mel	<i>Technical co-editor:</i> Michel Goossens
---	---

IF YOU NEED HELP (Contacts at CERN)

FRONT LINE Services

	Phone	Location	E-mail Address
Computing Help Desk	78888	513/R-052	Helpdesk@cern.ch
Opening hours: 9:00–12:30 and 14:00–17:30			
Administrative Applications (EDH, BHT, etc.)	79933	5/R-021	Ais.Support@cern.ch
Opening hours: 8.30–12:30 and 13:30–17:30			
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” (also known as the “UCO”—User Consultancy Office), is the proposed central point for help on all computer-related issues. Following is a non-exhaustive list of individual people responsible or assigned to a given service. If you do not find a pointer to the service you need, then contact the “Computing Help Desk” for submitting your request. For all the services listed below please use the generic electronic mail address whenever possible. These mailboxes will be read even when the usual specialist is absent. It is helpful to use relevant keywords in the subject field of your message. Contact specialists directly only for very urgent cases or for very general consultations.

Note that for 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 that page before submitting a question/problem.

USER SUPPORT

Service Definition

User Relations and Service Manager
Help Desk Manager
User Registration
Accounting Service
Computer BookShop

Name	E-mail Address
M. Marquina	helpdesk@cern.ch
R. Woolnough	helpdesk@cern.ch
M.C. Perler	user.registration@cern.ch
A. Koppanyi	helpdesk@cern.ch
J. Megies	cobs.service@cern.ch

CENTRAL COMPUTERS

Plus Services (Public Login Servers for interactive UNIX) P. Martucci pietro.martucci@cern.ch

CORE Services

– CSF (Simulation Cluster)
– SHIFT (Reconstruction & Analysis Clusters)
– PCSF (Windows NT simulation cluster)
– CORE Operations

V. Doré	csf.support@cern.ch
T. Cass	tony.cass@cern.ch
C. Boissat	pcsf-support@listbox.cern.ch
M. Vergari	mario.vergari@cern.ch

Work Group Servers
NAP and Parallel Applications
Direct Support for Physics
Central VMS Services
Oracle Databases
CERN-wide Print Service
Tape Allocations and Vault Space
Stager, tapes and TMS support
Central Data Recording (CDR)

T. Cass	tony.cass@cern.ch
E. McIntosh	eric.mcintosh@cern.ch
H. Renshall	harry.rensall@cern.ch
	vxcern.support@cern.ch
S. Santiago	oracle.support@cern.ch
J.L. Vosdey	printer.support@cern.ch
R. Minchin	richard.minchin@cern.ch
	tape.support@cern.ch
	cdr.support@cern.ch

DESKTOP COMPUTING

Service Definition	Name	E-mail Address
Desktop Infrastructure Services	A. Silverman	alan.silverman@cern.ch
Distributed File Services (incl. AFS)		afs.support@cern.ch
NICE (Networked PCs)	F. Hemmer	nice.support@cern.ch
MACintoshes Support	F. Ovet	mac.support@cern.ch
MACintoshes Sales	W. Hug	waltraud.hug@cern.ch
PC Desktop (Repairs and sales)	C. Ball	pc.admin@cern.ch
UNIX/LINUX Support	P. Defert	philippe.defert@cern.ch

Contact Accounts (by O/S)

Service Definition	E-mail Address
HP 700 s/w:	hp.support@cern.ch
SUN s/w:	sun.support@cern.ch
Solaris s/w:	solaris.support@cern.ch
RS/6000 s/w:	aix.support@cern.ch
Digital Unix s/w:	dunix.support@cern.ch
X terminals:	xterminal.support@cern.ch

SCIENTIFIC APPLICATIONS & SOFTWARE ENGINEERING

Service Definition	Name	E-mail Address
PaRC Service (Engineering Cluster)	P. Martucci	parc.support@cern.ch
CERN Program Library	L. Besson	heplib.support@cern.ch
PAW Support	O. Couet	paw.support@cern.ch
LHC++ and GEANT		heplib.support@cern.ch
Software Development Tools		sdt.support@cern.ch
Adobe FrameMaker		docsys.support@cern.ch

COMPUTING FOR ENGINEERING

All comments, problem reports, and suggestions can be reported via the “Feedback” Web form accessible from the “IT-CE Home page” (<http://wwwinfo.cern.ch/ce>, link “Feedback” on the left).

COMMUNICATIONS AND NETWORKS

A complete and up-to-date list of contacts is at the “Communications Systems Group” Home page at the URL: <http://network.cern.ch>

INTERNET APPLICATIONS

Service Definition	Name	E-mail Address
Mail support		mail.support@cern.ch
Web Office (support)	F. Fluckiger	www.support@cern.ch

COMPUTER SECURITY

Computer Security Information	J.Gamble	computer.security@cern.ch
-------------------------------	----------	---------------------------

OTHER COMPUTING NEWSLETTERS

COMPUTING FOR ENGINEERING: See the *CERN Computing Support for Engineering Newsletter*

URL: <http://wwwinfo.cern.ch/ce/newsletter/newsletter.html>

Distribution: M.Tate-Lavergne monique.tate-lavergne@cern.ch

Editor: L.Pregernig ludwig.pregernig@cern.ch

SERVICE URLs (Some Useful Web Pages)

TvScreen (Systems Status)	http://wwwcio.cern.ch/Others/TvScreen.html
Network Connection Request Form	http://network.cern.ch/register
Remote/Mobile Computing	http://network.cern.ch/mobile
Computing Helpdesk	http://consult.cern.ch/service/helpdesk
User Registration	http://wwwinfo.cern.ch/service/registration
COMputing Books Selling (COBS) Service	http://wwwinfo.cern.ch/service/bookshop
Accounting Service	http://wwwinfo.cern.ch/service/accounting/
Mail Services	http://consult.cern.ch/service/mail/
ListBox Mailing List Service	http://consult.cern.ch/service/listbox/
Netscape Service	http://consult.cern.ch/service/netscape/
FrameMaker Support	http://framemaker.cern.ch/
Text Processing related documentation	http://wwwinfo.cern.ch/asdoc/textproc.html
IT Home Page	http://wwwinfo.cern.ch/
User Support	http://wwwinfo.cern.ch/support
NICE Home Page	http://nicewww.cern.ch/
PC Support Service Page	http://pcwww.cern.ch/
PC Farms Homepage	http://wwwinfo.cern.ch/pdp/pc/
UNIX Workstation and X Terminals Support	http://wwwinfo.cern.ch/dis/ose/Support.html
Software Development Tools	http://www.cern.ch/PTT00LS/
Computing Support for Engineering	http://wwwinfo.cern.ch/ce
Joint Controls Project (LHC Experiments)	http://itcowww.cern.ch/ITC0/jcop/
Web Services	http://www.cern.ch/CERN/Divisions/IT/IA/services/website/
Web Office	http://www.cern.ch/WebOffice
Web Office / DWG	http://www.cern.ch/CERN/meetings/DWG/Members.html
Web Office / Guidelines	http://www.cern.ch/WebOffice/Guidelines/

SEMINARS AND PRESENTATIONS

Designation	Usual Location	Organizer(s)	Tel.	E-mail Address
Computing Colloquia	Auditorium / bldg 500	S. Jarp, IT	74944	sverre.jarp@cern.ch
Computing Seminars	IT Auditorium 31/3-005	G. Folger, IT	79778	computing.seminar@cern.ch
Technical Presentations	IT Auditorium 31/3-005	S. Cannon, IT	75036	stanley.cannon@cern.ch
Informatics Tutorials	IT Auditorium 31/3-005	M. Marquina, R. Ramos, IT		it-training.service@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.chSecretary: H. Meinhard/EP 1/1-028 76031 Helge.Meinhard@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 M. Draper	AT P. Heymans	IT A. Koppanyi	DSU M. Draper	FI M. Saitta
EST Ch. Delamare	PE D. Duret	EP A. Bjorkebo	PS J. Riche	SL M. Bornand
ST E. Sanchez-Corral	TH E. Gianolio	TIS G.R. Stevenson		

NICE Divisional RepresentativesSee the URL <http://nicewww.cern.ch/homepage/nicerepr.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" Home page, at the URL: <http://network.cern.ch>.

1. Announcements

This chapter only contains “punctual” (limited in time) or/and divisional announcements.

1.1 New Help Desk Service in Building 40

Nicole Crémel, Miguel Marquina and Roger Woolnough, IT/User Support

Following consultations with the ATLAS and CMS community the User Support team have established a “satellite” Helpdesk in Building 40. This service provides the B40 area and nearby with a personal contact to Help Desk Staff as well as a service for both book browsing and documentation supply.

As this service just started at the beginning of June the project is in its infancy and being moulded to suit the needs of the community. As a start the new desk at the entrance to Building 40 is staffed

Mondays, Wednesdays and Fridays 10:00-12:00

Long term goals of the Help Desk service are to extend these hours as needed and even open other such Desks providing the service is well used. We have taken advantage of the fact that the first month of opening co-incides with ATLAS and CMS weeks in an attempt to reach as many users as possible.

While we are still in negotiations with the Helpdesk service provider in order to accomodate these resources in the outsourcing contract, we will operate the new Desk with our existing resources; this means in practice that during the same three periods only one person (instead of the usual two) will man the Desk in building 513, therefore closing down the counter. Phone and email will still be regularly attended as usual, i.e.:

Phone: 78888

Email: Helpdesk@cern.ch

Web: <http://consult.cern.ch/service/helpdesk>

Do not hesitate to contact us (by email to User.Relations@cern.ch) for additional information or feedback regarding this matter.

1.2 Switching Off The Dial In Modem Service

Stanley Cannon, IT/CS

For a number of years IT division operated a modem dial in network which permitted (under certain conditions) worldwide access to the CERN computer network. For various reasons, mostly security and Y2K non compatibility, this access possibility will be definitively removed from service at the end of this year. In the interim the service has been reduced and the old modem lines have been disabled. The CERN recommended alternatives are to:

1. use ACB (<http://network.cern.ch/acb>)
2. log in via an ISP (Internet Service Provider) (<http://network.cern.ch/mobile>)

For those people who need time to migrate to an alternative method of access, please contact one of the following Email addresses: Stanley.Cannon@cern.ch, Netops@cern.ch or Helpdesk@cern.ch.

1.3 End of the CSF/NQS Service

Vincent Doré, IT/PDP

The CSF/NQS service started in 1991 will be stopped for ever the 30th of June 1999. Users should use instead the new CSF service running LSF for batch work, and based on more powerful processors that have been recuperated from other services when their hardware was

upgraded or replaced. More information on the new service can be found in CNL 233 at the URL (http://consult.cern.ch/cnl/233/art_csf.html). In case of problem, please contact CSF.Support@cern.ch.

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.

2.1 Invitation to Perform Y2K Testing

Sverre Jarp (CERN Y2K co-ordinator), IT/DI

Testing under UNIX

Introduction

A special AFS cell `y2k.cern.ch` has been established to allow service managers and users to test Y2K compliance.

In addition to AFS, the cluster consists of machines representing all the UNIX flavours in use at CERN (Aix, DUnix, HP-UX, IRIX, Linux, and Solaris).

More information can be obtained from the URL
<http://wwwinfo.cern.ch/pdp/bis/y2k/y2kplus.html>.

Testing schedule

The cluster will be set to 25 December 1999 on fixed days and then left running for three weeks. This gives people one week to prepare test programs in 1999 and two weeks to check the consequences of passing into year 2000. These fixed dates have been set as follows:

- 19 May 1999, date set to 25/12/99
(year 2000 starts on 26 May)
- 9 June 1999, date set to 25/12/99
(year 2000 starts on 16 June)
- 30 June 1999, date set to 25/12/99
(year 2000 starts on 7 July)

If more than these three sessions are needed an announcement will be made later.

Registration

The following Web page should be used for registration:
<http://wwwinfo.cern.ch/support/survey-info/y2kreg>

By registering your name and mail address on this page, you will be given a user-id on the first day of each test period. The userid will be valid for the corresponding period and will then be taken away for new usage. If you do not complete your tests in one cycle and wish to retain your account and the files in it, you must reissue the request via the registration web page.

Environment

The machines all run Y2K-compliant versions of the operating system. There is a reasonable amount of disk space in the y2k cell for home directories. Users can only access what is world-accessible in their real AFS directories in the CERN cell. File copying is therefore recommended (also to avoid that the Y2K-testing alters real production files or products).

Note that LSF will only be available as of the second session (starting 9 June). Tape access must be requested explicitly (on the Web page) as of this session.

The latest news, related to the test environment, will be posted on the Web pages mentioned above.

For PCs and Macs

Please refer to the IT Y2K page for further information:
http://nicewww.cern.ch/~sverre/Y2K_IT_Division.html

For ORACLE

Oracle database users who would like to test their applications under a Year 2000 system are invited to contact oracle.support@cern.ch.

2.2 Using a Single logonid for Multiple Groups or Experiments

Tony Cass, IT/PDP

Since April 21st people have been allowed to use a single logonid for all IT services even if they are working for more than one experiment or group. This means, for example, that it is no longer necessary to have accounts

JOAN on Shiftaleph for Aleph work, and
JOANZ5 on Rsplus for LHCb work.

Since April 21st it has been possible to register the JOAN account on Rsplus in group Z5.

How Does This Affect Me?

- If you have only one `logonid` and are working for a single CERN group or experiment then you are not affected.
- If you like having more than one `logonid` to keep your work for different groups or experiments separate then you are not affected. We will not force you to use a single `logonid`.
- If you are annoyed because your work for more than one group or experiment means you have to have more than one `logonid` then you can now ask your group administrator to let you use the `logonid` of your choice on all services—and this is likely to make it easier for you to access the Mail Servers. See below for more details.
- If you are a group administrator you need to understand the difference between Primary and Secondary accounts on a service and how these affect what a user can do—and what you can do for a user. See below for more details.

The account listings provided by `phone` and `xwho` have changed slightly to indicate the difference between Primary and Secondary accounts but most people won't need to worry about this.

What Does This Mean In Practice?

- People now have a Primary account on each service which associates the `logonid` to the group code normally used for work on that system. For example, the primary account for JOAN on `shiftableph` would be for group code XU and the primary account for JOAN on `Rsplus` would be for group Z5. These Primary accounts are the same as the old “ordinary” accounts—it is just that the group code associated with the Primary account can be different on each service.
- People may have a Secondary account on a service if they use that service for work associated with more than one group. JOAN, for example, could have a Secondary account on `Rsplus` for group code XU if she occasionally used `Rsplus` for Aleph work. To use the Secondary account the `newacct` command must be issued to set the correct group code, e.g.

```
newacct 'jdb$xu'
```

(The quotes are needed on Unix as the `$` character has a special meaning to the shell. The account code—especially the 3 letter UUU code—can be found from `xwho/phone` (`phone -A`) if you do not know it already.)

When submitting batch jobs with LSF, the `-P` option of `bsub` is needed, e.g.,

```
bsub -P xu < my_aleph_script
```

and the `newacct` command must also be issued from within the batch job.

- File space, in particular AFS space, is accounted only against the Primary account. It does not matter if the files contain Aleph or LHCb programs! Quota allocation can only be altered by the administrator for the Primary account—the account identified by a “\$” in `phone/xwho` output. (This change just extends the “home group” concept that always existed for AFS to other services and so the “home group” becomes the “Primary account”.)

How Does This Affect Group Administrators?

- Group Administrators can register an account on any service using an existing `logonid`. Many people will prefer this as it means they will keep easy access to their existing files and to the mail server as they change group. However, it will still be possible to create a new `logonid` for people who want to separate their work in different areas.
- Group Administrators can register a Secondary account on shared services (e.g. `Rsplus`) if the person concerned already has an account on that service in another group.
- Group Administrators will need to change Primary (and maybe Secondary) accounts as people move between experiments. For example, our Aleph physicist will have Primary accounts for group XU. As she starts work for LHCb some Z5 accounts (e.g. for AFS and maybe `Dxplus`) will initially be Secondary accounts. Once she is spending the bulk of her time working on LHCb these should be changed to Primary accounts with the Aleph accounts on these services becoming Secondary accounts.
- Group Administrators must realise that blocking access to a service is at the **USER** and not at the **ACCOUNT** level. This means that if **ANY** account is blocked (whether Primary or Secondary) then the user concerned will not be able to use that service at all. Please look to see if any other accounts, whether Primary or Secondary, exist for another group before blocking a user account. It may be better to change the group code for the Primary account or to switch the Primary and Secondary accounts and then delete (not block) the new Secondary account.

2.3 New STK 9840 Tape Units Enter Service

Charles Curran, IT/PDP

A new Powderhorn silo...

Visitors to the Computer Centre will have noticed that the existing Storage Tek tape library has recently been extended. It is now configured as 4 Powderhorn 9310 silos, each capable of holding approximately 6,000 cartridges. All are interconnected.

...new tape units also...

The most recently added silo is equipped with 10 of STK's newest tape drive, the 9840. This unit is considerably cheaper than the *Redwoods* already installed, but offers a much higher performance than what are sometimes called "commodity" tape units (in which category we include the DLT 2000 unit, for example).

Read and write speed is 10 MBytes/sec, and the cartridge used by the unit can hold up to 20 GBytes of data. These are the native speed and native capacity, which ignore any effects of the data compression which the unit also offers. Recording is linear serpentine, not helical, over 288 tracks. See STK's web pages at URL: <http://www.storagetek.com/StorageTek/hardware/tape/9840/> for details.

See CERN's web pages at URL:

<http://wwwinfo.cern.ch/pdp/vm/picture/stk9840.gif> for a view of CERN's 9840s.

...which use new media...

While these are already quite impressive performances, the 9840 has another novel feature. The cartridge or *cassette* ("cassette" describes its design more clearly) contains tape which is wound onto two reels contained within the "3480" form factor. It thus requires no complicated threading into the unit. The cassette, once presented to the unit by the Powderhorn robotics, is loaded and initialised ready for operation in only 4 seconds, and can seek to any data on the tape in a little

as 8 seconds. This is much faster than either the STK *Redwood* or the IBM 3590, and makes this media and unit particularly suited to accessing relatively scattered files on one or more cassettes rapidly. There are 300 such cassettes in the library, which carry pre-printed stickers (R01001-R01300, with media identifier "R"). See CERN's web pages at URLs:

<http://wwwinfo.cern.ch/pdp/vm/picture/9840cart1.gif>

and for views of a 9840 cassette, see the URL

<http://wwwinfo.cern.ch/pdp/vm/picture/9840cart2.gif>

What will they be used for?

The characteristics noted above may well be quite attractive to many users, but the cassette itself is *more expensive per GByte* than the *Redwood*, so we do not expect it to be used for an application such as bulk recording of raw experimental data. However, there are some cartridges at CERN that are accessed very often, which are used rather regularly in "append" mode for journalling or in "append" mode for some relatively low-volume Central Data Recording applications.

We anticipate that its characteristics will allow us greater flexibility in our storage system plans, offering a device that is cheaper than disk, but much faster in providing access to data than any of our existing tape units. Systems such as HPSS can exploit devices of different characteristics without explicit user intervention rather easily by using an appropriate "Class of Service"; improvements to the current SHIFT software will probably do the same. It should also be simpler to offer a service which users see as "high reliability" with such a unit: considerable effort is still required to provide such a service with *Redwoods*.

We will try to contact users with appropriate usage patterns to see if they would like to use this new offering. Equally, if you are interested, you can send an e-mail to tape.support@cern.ch, and we will be happy to discuss what use you might wish to make of the 9840.

2.4 Connecting to CERN from Home: ACB Service over ISDN

John Ogilvie, IT/CS

Overview of the ACB ISDN Service

ISDN (Integrated Services Digital Network) is an extension to the public telephone network, which guarantees higher speed (64 or 128kbps) network access.

An initial ACB (Automatic Call Back) ISDN service, based on a small ISDN router connected to a home computer (PC, Mac or Linux) via an Ethernet card is available. It currently offers only Internet access (direct access to NICE and Appletalk files is under investigation). *To use ISDN your home must be cabled* for digital access to your local Telecom supplier.

The *ACB ISDN* service is offered in addition to the existing *ACB modem* service. The *ACB modem* service will continue to be operated and improved (56kbps modem access is planned for later this year).

Choosing between the ISDN and Modem ACB services

ISDN Advantages:

- guaranteed throughput at 64 or 128 kbps (see Note 1 below);

- low latency—giving good interactive response time;
- fast connection time;
- can make phone calls while using the network.

Modem Advantage:

- no extra monthly rental or installation costs (see Note 2 below)

Note 1 ISDN provides 2 guaranteed digital channels at 64kbps, whereas actual throughput achieved by 56kbps modems is typically around 40 kbps (depends on the line quality). In addition, 56kbps modems are asymmetric with only 33.6kbps in one direction (home to CERN).

Note 2 Installing ISDN in your home costs significantly more than an analogue telephone line. ISDN monthly charges are roughly equivalent to the rental of 2 analogue lines.

For more complete information, please see the ACB Web pages: <http://network.cern.ch/acb>

3. Desktop Computing

It includes CERN desktop environment centrally provided on UNIX and PCs (NICE95 and NICE-NT), desktop software strategy, etc.

3.1 *Springer* UNIX Print Server Removal

Jean-Louis Vosdey, IT/CIO (Printer.Support@cern.ch)

For a long time now *Springer* has been the major Print Server in the computer centre. It continues however to run a very old version of Ultrix that is not Y2K compliant. For this reason every effort is being made to move printers away from this particular server and onto the new PRINT service composed of a series of LINUX PC's which are Y2K compliant. Following the current schedule this move should be completed by September of this year and *Springer* taken out of service.

As a result of these changes the following has become necessary:

xprint as used by UNIX users has been rewritten.

- Supported platforms can use the **xprint** command in PATH /usr/local/bin/xprint from the ASIS software repository.
- If your UNIX environment is not standard you have to download **xprint** via anonymous ftp on asisftp.cern.ch.
- Explanations and answers to your questions can be found at the URL address:
<http://wwwcio.cern.ch/Printers/cnl.html>.

From now on, the new PRINTER WIZARD enables users to print from W95, W98 and WNT PC platforms, whether they are connected to NICE or not.

- Additionally, there is now only one print queue from both platforms (PC and UNIX).
- The package and explanation can be found at the URL address:
<http://printdb.cern.ch/cernprinting/>.

Furthermore, the gateway (PCCI004) between windows NT users and *Springer* has become redundant and will be also taken out of service at the end of September. We strongly recommend users to use PRINTER WIZARD on windows NT.

On VMS, work is in progress on a replacement for WPRINT (the VMS equivalent of **xprint**) which will bypass *Springer*.

Note that an update of any modifications or new features, where printing is concerned, will appear in future issues of the Cern Computer NewsLetter.

3.2 New Print Package on NICE

Jean-Louis Vosdey, IT/CIO (Printer.Support@cern.ch)

A new CERN Printing Package has been developed in order to improve printing reliability and control as well as to enable non-NICE PCs to print on CERN's network printers. It was tested during the months of April and May by about 100 users.

As announced in the last Desktop Forum, this new package has gradually replaced the current NICE Printer Wizard during the month of June, as explained in the following plan:

- June 8 : IT and PS divisions
- June 22 : AC, AS, DG, DSU, EST, FI, LHC divisions
- June 24 : EP, PE, SPL, TH, TIS divisions
- June 28 : SL, ST divisions

For users who want to use the new Printing Package before the deployment date, it can be installed by running the following program (only once):

```
z:\p32\cernprinting\package\printrun.exe
```

Important: Newly installed printers (after April 99) can only be accessed from the new Printing Package.

This new package is completely transparent for the NICE95 and NICENT users. The Printer Wizard icons in the 'Control Panel' and in the 'Start Menu->More Applications->System Configuration' will automatically refer to the new version. You will still be able to access the Printer Wizard via these icons, but a new Task-bar Icon in the 'Icon Tray' (right side of the task-bar, next to the clock) will also provide access to the Printer Wizard.

For more details about this icon see the URL

<http://printdb.cern.ch/cernprinting/client.html#PrntTray>

It is highly recommended to use this new icon, which will give you access to the following (additional) features by selecting it with the right mouse button:

1. **Printer Wizard:** selecting this option will launch the new Printer Wizard (identical user interface as the previous version). See also <http://printdb.cern.ch/cernprinting/client.html#PrintWiz>
2. **Job Information:** a light-weight program has been created to replace the job info window in the Printer Wizard. This new tool starts up very quickly and enables the user to control the printer jobs and settings without starting the more heavy Printer Wizard. See also <http://printdb.cern.ch/cernprinting/client.html#JobInfo>

3. **Access to the LPR Service window:** the new version of the LPRServ background program includes the following improvements:

- The window only pops up in case of a printing error. In normal conditions, an icon appears on task bar during print job transfer.
- The window returning a history of all printed jobs can be activated by selecting the 'Show LPR Client Service' option.

See also the URL <http://printdb.cern.ch/cernprinting/client.html#LPRServ>

More information about this new CERN Client Printing Package is available at the URL <http://printdb.cern.ch/cernprinting>.

3.3 Windows NT *Service Pack 4* Installation on NICE NT Systems

Alberto Di Meglio, IT/DIS

Overview

Service Pack 4 for Windows NT 4 was released at the end of 1998 and contains various updates and bug fixes for Windows NT.

The most important improvements introduced by SP4 can be summarised as follows:

- Certified Y2K compliance: after application of SP4 and the post-SP4 Hot-Fixes, the system is rated COMPLIANT by *Microsoft*.
For more information on Y2K compliance issues of Windows NT 4 check the URLs at:
http://microsoft.com/technet/year2k/product/user_view68399EN.htm
(Windows NT 4 Workstation)
http://microsoft.com/technet/year2k/product/user_view67961EN.htm
(Windows NT 4 Server)
- System Environment Block problem fix (currently bugging many *VisualStudio* users): if the user environment is large (long path, many environment variables defined), the *winlogon* service may corrupt the System Environment Block and the workstation may crash with STOP error 0xC000021A when shutting down Windows or using Task Manager. This has been seen to happen on computers where *VisualStudio* and other large development programs had been installed.
- Euro symbol support.
- Additional support for *NetWare File* and *Print Sharing*.

- Updated *Time Zones* definitions
- Security configuration and management tools (server version only, it is not included in the standard NICE installation)
- Automatic installation of AGP graphics card during setup: Windows NT does not have direct support for the AGP bus. Therefore when an AGP graphics card is present in the system, it is not automatically configured during installation. Support for AGP is introduced by SP3 and SP4, which have to be installed before the AGP card drivers. The installation set will now be patched to configure AGP cards automatically
- Improved access to IDE disks with more than 8GB space during setup: SP4 contains a new IDE bus driver that allows to partition correctly IDE disks larger than 8GB during setup.

For more detailed information about the fixes contained in SP4, check the URLs at:

<http://www.microsoft.com/ntserver/nts/exec/overview/NT4SP4whatnew.asp>

(overview of all new features and fixes)

<http://support.microsoft.com/support/kb/articles/q194/3/03.asp>

(detailed technical list of all fixes)

Installation instructions of the NICE NT SP4 package

SP4 will be deployed in the NICE NT environment during the months of June and July one Division at a time.

The installation of SP4 on NICE is automatically started when a user logs in a computer that has not yet received the update.

In order to install SP4, it is necessary to be logged under an account with local administrator privileges. If the account used does not have local administrator privileges, a dialog box will prompt you to log in under a valid account to start the procedure and the installation program will then stop. It is not possible to start accidentally the installation while logged in as a non privileged user. You can have a look in the NICE NT FAQ page at <http://nicentinfo.web.cern.ch> for more information about accounts with local administrator privileges (*I want my NICE account to have local administrator privileges. How do I do it?*, see <http://nicentinfo.web.cern.ch/nicentinfo/thenice1.htm#Q>).

The installation consists of one or more steps, depending on the components that you will choose to install.

The first time that you log in under a privileged account, you will see a dialog with four buttons, one for each option.

Option 1: "Yes, proceed with the installation"

If you choose the first option, the installation of SP4 will immediately proceed and will install the basic components required to update your system. You need to have about 35 MB of free disk space to install SP4. You will see a dialog showing you the installation progress. Do not use the computer during this phase. At the end of the installation, the computer will automatically reboot. At this point you can log in with your normal account.

Option 2: "Maybe, but I want more information"

Select this option if you want to choose which components to install. A dialog in French or English will guide you through a series of steps. For the moment the additional components you can install are *Internet Explorer 4 Service Pack 1* and the updated Data Access components (v.2.0) for working with databases. You need to have up to 80 MB of free disk space for this option.

This installation is divided in three phases:

1. During the first phase, the system updates are installed. You will see a dialog showing you the installation progress. Do not use the computer during this phase. At the end of the installation, the computer will ask you to confirm before rebooting
2. After the reboot, log in again as local administrator to start the second phase. During this phase, IE4 SP1 and the Data Access components will be installed. The computer will then reboot automatically at the end of the procedure without asking for a confirmation.
3. After the second reboot, log in again as local administrator to complete the installation of IE4 SP1. At the end of this phase, you can use your computer as normal and you can log out and log back in with your normal account.

Option 3: "Not now. Remind me again next time"

Select this option if you want to postpone the installation to the next login as local administrator.

Option 4: "Not now. I'll do it myself later"

Select this option if you do not want to install SP4 and do not want to be reminded anymore.

Manual installation

Service Pack 4 can be installed manually at any time from the Windows Menu by going to

Start - More Applications -
System Configuration - PC Administration -
Windows NT 4 SP4 Setup

The manual installation does not check for the user Division and can be used to install SP4 even before your scheduled turn.

Re-installation of SP4

If you want to reinstall SP4, you need to select the first option. IE4 and the Data Components cannot be re-installed using this procedure. If you need to reinstall *Internet Explorer 4.0 SP1* you have to use the procedure from:

Start - More Applications -
World Wide Web - Install Internet Explorer,
Netmeeting and Chat

Known bugs and limitations - Please read carefully before installing SP4

- On some computers during the installation of the basic components (Option 1 or Option 2 - Phase 1), we have seen an error message stating that the file IEXPLORE.EXE cannot be found. This is not true, but it depends on the fact that a copy of the file has been found on one of the NICE servers and the setup program is trying to overwrite that file. Just press the Skip button to proceed.

IMPORTANT: if you see this error message, **DO NOT** reset the computer or you will not be able to connect to the CERN NT Domain anymore. Just press the Skip button.

- On some computers, during the third phase of the complete installation (Option 2), we have seen one or more error messages stating that some *Java* class could not be registered. Press on the OK button to proceed. The classes will be registered again the first time you use them.

- If you receive the prompt to install SP4 while installing or reinstalling a NICENT system, postpone the installation until NICE has completely finished updating your system. This is necessary to be sure that all files to be updated are already

present on your disk.

For additional questions and comments, contact the author (Email Alberto.Di.Meglio@cern.ch).

3.4 Y2K Compliance of UNIX Systems

Lionel Cons, IT/DIS

The latest certified version of all the supported UNIX systems at CERN is, according to their respective vendors, Y2K-compliant. That is:

- AIX 4.3.2
- Digital UNIX 4.0d
- HP-UX 10.20 ACE
- Red Hat Linux 5.1
- Solaris 2.6

We are anyway testing these systems in real conditions in the “Y2K cluster” (see the article *Invitation to per-*

form Y2K testing under Unix on page 2 of this CNL). Note also, that some vendors may release Y2K related patches later this year so you should keep an eye on the URL:

http://nicewww.cern.ch/~favrot/year_2000.htm.

We recommend to upgrade your UNIX machines to the latest certified level as soon as possible to avoid problems when approaching the end of the year. Note however that we will very likely certify a new version of Linux later this year (this was discussed at the next CLUG meeting, in June) so, for Linux only, you may want to wait before upgrading.

3.5 Operating System Certification - Present List

Philippe Defert, IT/DIS

As stated in CNL 224 here is the list of various versions of UNIX which we certify as working correctly with the CERN-recommended libraries, tools and utilities as of March, 1999:

- Linux Redhat 4.1/4.2 - certified and frozen
- Linux Redhat 5.1 - certified and recommended
- Solaris 2.5.1 - certified and frozen
- Solaris 2.6 - certified and recommended
- HP-UX 10.20 - certified and frozen
- HP-UX 10.20 ACE - certified and recommended

- Digital UNIX 4.0.b - certified and frozen
- Digital UNIX 4.0.d - certified and recommended
- AIX/6000 4.1.5 - certified and frozen
- AIX/6000 4.3.2 - certified and recommended

For those interested in the certification process, it is described in the Web page:

<http://wsspinfo.cern.ch/file/certified.html>

The latest news on the certification procedure is always published in the newsgroups associated with each UNIX flavour (for example news:cern.linux, etc.) and summarised in the Web page listed above and in each subsequent CNL.

3.6 Operating Systems Availability on ASIS

Philippe Defert, IT/DIS

The status with the ASIS Application Software repository is presently as follows:

Supported platforms:

- **i386_redhat51:** Linux RedHat 5.1
- **i386_linux2:** Linux RedHat 4.1 (frozen)
- **sun4x_56:** SUN Solaris 2.6
- **sun4x_55:** SUN Solaris 2.5.1 (frozen)
- **hp_ux102:** HP-UX 10.20
- **alpha_dux40:** DEC OSF 4.0b
- **rs_aix41:** IBM AIX 4.1.5 (frozen)
- **rs_aix43:** IBM AIX 4.3.2

Unsupported platforms:

- **sgi_64:** SGI Irix 6.4
- **sun4c_411** (SUNOS 4.1.3) (frozen since 1997)
- **hp700_ux90** (HP-UX 9.05) (frozen since 1997)
- **alpha_osf32** (DEC OSF 3.2c) (frozen since 1997)

A platform is declared *frozen* when all applications are kept in the repository but no new package or version of a package is ever introduced in the repository for this architecture.

A platform is declared *partial* when only a part of the applications is generated for this particular O.S.. Getting access to ASIS applications is to be done with the appropriate tools which will give you some application of this particular platform and the rest from the closest O.S., called the *fail-over*.

4. Internet Services

Everything related to Internet in a large sense: i.e., all technical Web issues, electronic mail, Internet news, home access to CERN facilities, video conferencing and multi-media, network security, etc. Incidentally please note that all Web activity at CERN is governed by the Web Policy Group (WPG) and the Divisional Webmasters Group (DWG). The Web Office coordinates the day-to-day tasks. The DWG provides Guidelines for Web sites at CERN; when in doubt, consult your Division's responsible member.

4.1 Configuring an Internet Access Provider

Denise Heagerty, IT/CS

As explained in CNL 232 (*Remote/Mobile Computing at CERN*, page 3, see also the URL

http://consult.cern.ch/cnls/232/art_mobcomp.html

or the Remote/Mobile Computing Home Page <http://network.cern.ch/mobile>) is the focal point for user requirements regarding this issue. This page has been up-

dated recently to provide help for *Connecting to CERN whilst on travel*, and more specifically:

- connecting to CERN via Internet Access Providers
- configuring both an Internet Access Provider and CERN ACB

4.2 Computer Networking is Getting Faster and Better Organised

Jean-Michel Jouanigot, IT/CS

The structured cabling project as defined in 1994 was completed by the end of 1998. The main goal of this project was to provide reliable 10 Mbits per second (Mbps) network connections to the desks in all of the major buildings. Its companion, the routing project, whose goal was to restructure and stabilize the overall network, is also complete. In all of the major buildings, structured cabling and routing have provided reliable and stable network services to almost all the desks, and to the Computer Centre. The backbone itself, defined at the beginning of this complex project, was based upon switched FDDI 100 Mbps rings, interconnected by a series of Gigaswitches. On each ring a series of routers are connected, providing shared 10 Mbps Ethernet to most of the desks, or FDDI to some work groups.

At some places structured cabling could not be introduced: in these the remaining coaxial infrastructure is being restructured and then routed, thus creating what we can call "structured coax". In particular, the experimental halls have all been restructured, such that each experiment now gets its own dedicated Ethernet segment(s). It is expected that this restructuring and routing will be completed by the beginning of next year, at which time the old network (in which all TCP/IP network numbers are of the form 128.141.x.y) will be completely removed.

Since the definition of the project some new technologies have arrived on the market. The IT/CS group has been closely looking at these emerging technologies, and has evaluated some of them, in particular *Fast Ethernet* and *ATM*. Since 1997 it has become clear that *Fast*

Ethernet would become the de facto standard for 100 Mbps connections. *Fast Ethernet* is exactly like standard Ethernet, but running 10 times faster. However, it can also run in "Full duplex" between two systems, with the advantage that there are no more collisions on the medium. It is therefore possible to get more than 11 Megabytes per second (MBps) between two systems in each direction at the same time (a value which is easily attainable with modern systems). *Fast Ethernet* can only run on Fibre or on the structured cabling wiring and plugs: it cannot run on the old coaxial cables. The first *Fast Ethernet* switch was installed in August 1997 in the Computer Centre, and the first experiment to use this technology was NA57 in their Data Acquisition System (DAQ) at the end of 1997. It was clear that deploying *Fast Ethernet* connections would not be possible if the CERN backbone itself was not able to handle multiple simultaneous connections. *Gigabit Ethernet* was available just in time.

Gigabit Ethernet is exactly like *Fast Ethernet*, but running 10 times faster (100 times faster than Ethernet). It can only run Full duplex between two systems, meaning that the available bandwidth is more than 120 MBps in each direction at the same time. It is currently only available on fibre, but it is expected the structured cabling and plugs will very soon be able to carry that speed! The first *Gigabit Ethernet* experiments were carried at CERN in late 1997, and used in production for the first time for the NA48 Central Data Recording in January 1998 between the North Area (Building 918 in Preessin) and the CS2 in the Computer Centre. The

switches used in this experiment were two FDDI/*Fast Ethernet* switches interconnected with *Gigabit Ethernet* over a 7 km link (which is almost twice the accepted maximum length defined by the standard). Four systems in the NA48 DAQ were connected with FDDI to one of the switches, whilst the CS2 was connected to the system using four *Fast Ethernet* connections. The maximum effective bandwidth achieved (real data, real software), was over 35 Mbytes per second (30% of the Gigabit link), which was more than the experiment's requirement. This setup ran successfully until the beginning of 1999, at which time NA48 was integrated into the general CDR Gigabit facility.

Since then, progressively, the *Gigabit Ethernet* infrastructure has been deployed, strongly pushed by the needs of the PC farms used in physics data analysis. Some of these farms include more than 50 PCs interconnected with *Fast Ethernet*, the disk servers and tape servers being connected with *Gigabit Ethernet* to the switching system. More than 10 such farms have been installed in the last 18 months in the Computer Centre or in the physics DAQs.

Since the number of systems using *Fast Ethernet* or *Gigabit Ethernet* was rapidly increasing, the introduction of a new generation of routers, able to handle Gigabit speeds, was necessary. The first two such routers were introduced in the Computer Centre in February 1999: all central PC farms, as well as most of the NICE servers,

are now connected to them. The migration to these new routers of the various other servers in the computer centre is planned this year. At the same time, the central FDDI Gigaswitch was replaced by new FDDI/*Gigabit Ethernet* switches so that the old FDDI backbone now has a fast connection to the new one.

The IT/CS group has started the deployment of the new generation *Campus backbone* based upon Gigabit routers interconnected with Gigabit links. The very first places where this new backbone was deployed are the experimental halls. In Preessin the North Hall (around building 887) has CMS, NA49, NA57, NA48 and COM-PASS, all of which have their Data Acquisition systems connected with *Gigabit Ethernet*. The East Hall (Building 157) and the West Area (building 180) also now are on the new backbone. This overall network setup was shown to permit more than 110 Mbytes per second between the NA48 DAQ in Preessin and a series of servers on *Gigabit Ethernet* in the Computer Centre.

This new backbone will in future be deployed to other places, where required. It will provide high bandwidth to the structured cabling star-points, making possible the introduction of *Fast Ethernet* on the desk where necessary. It is currently planned to equip more than 10 star-points with this new technology by the end of this year.

For more information on this new project, please consult the URL <http://network.cern.ch/> (section "Projects").

4.3 How (Not) to Block the Computer Networking and... ... (Not) Then Get Blocked by Network Operations!

Mike Gerard, IT/CS

The accompanying article (page ~11) on the improvements in computer networking in CERN has indicated the considerable effort which is going into restructuring, creating a high performance backbone and offering the possibility of central data recording at very high rates of transfer. However, limitations of money and manpower have led to decisions on priorities such that most desktop equipment (PC, Macintosh, X-terminal or Unix workstation) is still sharing a 10 Megabit/second Ethernet segment with around 40 other systems. Given the power of modern processors, almost any desktop system can effectively monopolise such a segment for long periods. When we discover such cases, normally as a result of other users complaining, we often have no option but to block the offending system in order to protect the majority of users. However, we would much rather that users try to avoid any activity which could cause such situations.

Sometimes the problem is caused by a workstation trying to process large amounts of data which are actually held elsewhere: often an AFS or NFS mounted file held

centrally somewhere. Since the workstation may have powerful CPU capacity the processing may be rapid, thus causing high data transfer. Our objective is rather to move this processing capacity to the data, hence the increasing number of PC processing farms in the Computer Centre.

The most frequent problem, however, is where an executing program is not running in the local workstation, but is exporting information to go onto the screen (and loudspeakers, perhaps) of the workstation. One example of this (which we now try to ban) is the use of a screen locking program (xlock) used to lock an X-terminal when a user is absent, and which sends constantly changing patterns to the locked screen. The answer is: use a simple screen lock which uses a blank screen!

More difficult (even impossible) to avoid is the use of the *Netscape* Web browser to view a dynamic page (perhaps displaying dynamic graphics files, perhaps running a Java applet). If the browser is executing remotely

from the X-terminal (or a workstation running as an X-station, such as a PC running Exceed), then even a harmless-looking object (such as a rotating globe in a corner) can take a large fraction of the available network bandwidth. And that remains true even if the rotating globe is not visible on the screen, or even if the whole *Netscape* window is reduced to an icon: it can, however, be stopped by turning off animation (in the *Netscape* “View” Menu). Therefore, for those people who can only run *Netscape* remotely, we ask that they be aware of these possible network monopoly situations and avoid them.

A number of persons who have personal workstations capable of running *Netscape* locally are nonetheless choosing to run it remotely, often on the central PLUS servers. Arguments for doing this can be found (central bookmarks, modern version of *Netscape*, lack of performance of local workstation). However, there are obvious counter arguments (put bookmarks onto a central AFS, NFS or Novell server!). The plain truth is that from a networking point of view it is *almost always* more efficient for our networking if *Netscape* is run from the local

Operating System (i.e., your personal desktop machine).

Netscape can also be used together with various plug-in programs for receiving on-line audio/video. As long as this is in some way work-oriented there is normally no problem. However, if (as is technically possible) a user retains a connection to an external audio/video source for a long time then the data transfer becomes significant. Clearly, although the objective of CERN providing computing and networking facilities is to enable people to do CERN work, no-one is going to object to occasional casual surfing (e.g., of news headlines on www.cnn.com or similar sites). However, staying for significant periods of time on sites offering audio/video feeds unrelated to CERN can have an adverse effect on other users.

The bottom line is that for many people at CERN the potential for overloading the local network segment is a real one. Therefore, please try to recognise and avoid such situations before we are obliged, as a result of complaints from other users, to “pull out the communications plug” on your desktop system.

4.4 CERN’s Connections to European and Global Internet Improved

Olivier Martin, IT/CS

Significant improvements have been achieved this year with the deployment of TEN-155, the new 155 Mbps pan-European backbone interconnecting the National Research Networks (NRN). CERN is sharing a 155 Mbps access point with SWITCH, the Swiss NRN and has a guaranteed bandwidth of 40 Mbps, 5 times more than what we had in 1998 on the TEN-34 backbone.

Another major milestone was achieved in April with the re-deployment of our dedicated transatlantic line from MCI/Perryman to C&W/Chicago, whilst simultaneously increasing the capacity from 4Mbps to 12Mbps initially then 20Mbps by August 1999. Very soon, CERN will connect from its Point of Presence in Chicago to the Intercontinental Internet Exchange Point in Chicago, sponsored by the US National Science Foundation and known as STARTAP (Science Technology and Research Transit Access Point). This will give us direct access

to a number of US, Canadian and Asia-Pacific networks such as ESnet, vBNS, Internet2/Abilene, Canarie and Transpac.

Our transatlantic connectivity is organized through a consortium led by CERN, with IN2P3, WHO, US and Canada HEP communities through Caltech and Carleton University.

Thanks to exceptional circumstances causing the price/bandwidth ratio to drop very significantly, the above upgrades could be achieved at approximately constant budget levels. The result is that we now start to have adequate capacity for all normal CERN requirements, in particular for exporting sizeable amounts of physics data all over the world, in a coordinated and controlled manner. Further improvements are expected next year.

5. Scientific Applications and Software Engineering

This includes descriptions of:

- *scientific applications and libraries from CERNLIB and LHC++, mainly used by the Physicists,*
- *engineering applications like the ones on PaRC or those associated to online controls,*
- *software development environment.*

5.1 Garfield Version 6.29

Rob Veenhof, GSI (Darmstadt, Germany)

Garfield version 6.29 is a maintenance release with enhancements in the areas of field map interfaces, of signal calculations, and of arrival time distributions for electrodes other than wires.

CELL section:

An interface for Tosca generated field maps with parallelepipeds as elements has been written in collaboration with Guido Maria Urciuoli (GuidoMaria.Urciuoli@iss.infn.it) at the INFN Gruppo Collegata Sanità.

Field maps with both an axial periodicity around an axis and a translation periodicity along the same axis, can now be processed.

Elimination of the excluded parts of the background in Maxwell Field Simulator 3D field maps has been improved.

The WIRE-MARKERS option has been made available in the field, drift and signal sections.

Occasionally, plotting the materials present in field map resulted in the screen being coloured uniformly. A new attempt has been made to avoid this.

MAGNETIC-FIELD section:

Magnetic fields computed by finite element programs can now be entered via the FIELD-MAP statement in the cell section.

OPTIMISATION section:

No changes.

FIELD section:

No changes.

GAS section:

Tracing Heed generated delta electrons in external electric and magnetic fields would occasionally stop program

execution when the delta left the Heed tracking volume. This has been corrected.

DRIFT section:

Two new clustering models have been added to the TRACK statement. They generate points along the track that are spaced according to the flux of the electric field.

The ARRIVAL-TIME-DISTRIBUTION command has been extended to deal with other electrodes than wires. The command now also operates in the y-z and z-y planes. Existing input files which use this command need to be modified.

Track interpolation has been improved: options have been added to INTEGRATION-PARAMETERS which give better control over the interpolation order and the way delta electrons are dealt with.

Drifting from the intersection between the viewing plane and selected solids has been added.

SIGNAL section:

Signal calculation in other electrodes than wires (planes, tube, solids) has been enabled. Electrons arriving on a plane, tube or solid now produce a signal that is classified as direct.

Graphics:

No changes.

Histograms, Matrices, Formulae and Calls:

No changes.

Datasets and input / output:

No changes.

Compilation, batch:

No changes.

Documentation

Currently, the most accurate information can be found in the help pages at the URL

<http://consult.cern.ch/writeup/garfield/help>.
The description of the command line is at
<http://consult.cern.ch/writeup/garfield/command.html>.

5.2 Software Development Tools Service

Eric Poinsignon, IT/IPT (sdt.support@cern.ch)

In IPT Group we actively help software projects in CERN experiments to identify, evaluate, purchase and use industrial software development tools. Network floating licences are used as much as possible, so that the tools can be run remotely. We can also, if you have an AFS account, give you access to our Solaris server in case the binary for your platform is not available. You can find the complete information on what we have and how to use it at URL:

<http://www.cern.ch/PTTOOLS/>

Change of E-mails

Our two “service e-mail addresses” are now in production. To contact the SDT team send mail to one of the following emails depending on your query:

Sdt.Support@cern.ch To get access to a tool in production, for help and advice on any tool in production (see below for **FrameMaker**, **IXgen** and templates), and general inquiries on the service. This replaces the old swtool@ptsun00.cern.ch address.

Docsys.Support@cern.ch For help and advice on using **FrameMaker**, **IXgen** and templates. This replaces the old docsys@ptsun00.cern.ch address.

Below is the complete list of our tools. For the latest information watch our News page which is updated as things happen:

<http://www.cern.ch/PTTOOLS/SoftNew.html>

Analysis and Design

- **ADAMO** (IPT product): ER design and programming system.
- **Rose**: OMT/Booch/UML analysis and design tool with C++/Java/Ada code generation.
- **StP**: OMT/Booch/UML analysis and design with C++ code generation.

Modelling and Simulation

- **Modsim**: Discrete event modelling and simulation tool.

Programming and Debugging

- **SNiFF+**: Integrated development environment for C, C++, Java, Fortran and CORBA IDL

Graphical User Interface Builder

- **X-Designer**: Interactive X-platform GUI builder with OSF/Motif toolkit widgets, C, JAVA and Windows code generation.

Quality, Metrics and Testing

- **(changed)** **Logiscope Tool Kit**: Source code analyser and test coverage for C/C++/Fortran.
- **Look!**: C++ dynamic visualization system.
- **Insure++**: Comprehensive tool for error detection in C/C++ programs.

Documentation

- **FrameMaker**: Complete and integrated publishing tool.
- **IXgen**: Indexing and marker management tools for FrameMaker.
- **(new)** **SoDA**: Software Documentation Automation.
- **(changed)** **VisualThought**: Multipurpose drawing, diagramming and flowcharting tool.
- **(new)** **SDLT** (IPT product): Software Documentation Layout Templates, for FrameMaker.
- **User Guide Templates** (IPT product): For FrameMaker.

6. The Learning Zone

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

6.1 A Tutorial on Character Code Issues

Jukka Korpela, Helsinki University of Technology (HUT), Computing Centre, Finland

Note: this article is a shortened version of a more detailed document available on the Web at the URL <http://www.hut.fi/u/jkorpela/chars.html>. As it contains many references which are links to other Web pages it is preferable to browse it from the Web.

This article tries to clarify the concepts of *character repertoire*, *character code*, and *character encoding* especially in the Internet context. It specifically avoids the term *character set*, which is confusingly used to denote repertoire or code or encoding. ASCII, the ISO 8859 family (ISO Latin, especially ISO Latin 1 and MIME are used as examples. The document in itself does not contain solutions to practical problems with character codes; rather, it gives background information needed for understanding what solutions there might be, what the different solutions do—and what is really the problem in the first place.

The basics

In computers and in data transmission between them, i.e. in digital data processing and transfer, data is internally presented as octets, as a rule. An octet is a small unit of data with a numerical value between 0 and 255, inclusively. The numerical values are presented in the normal (decimal) notation here, but notice that other presentations are used too, especially octal (base 8) or hexadecimal (base 16) notation. Octets are often called bytes, but in principle, octet is a more definite concept than byte. Internally, octets consist of eight bits (hence the name, from Latin octo 'eight'), but we need not go into bit level here.

Different conventions can be established as regards to how an octet or a sequence of octets presents some data. For instance, four consecutive octets often form a unit which presents a real number according to a specific standard. We are here interested in the presentation of character data (or string data; a string is a sequence of characters) only.

In the simplest case, which is still widely used, one octet corresponds to one character according to some mapping table (encoding). Naturally, this allows at most 256 different characters being represented. There are several different encodings, such as the well-known ASCII encoding and the ISO Latin family of encodings. The correct interpretation and processing of character data of course requires knowledge about the encoding used. For HTML documents, such information should be sent by the Web server along with the document itself, using

so-called HTTP headers.

Previously ASCII encoding was usually assumed by default (and it is still widely used). Nowadays ISO Latin 1, which can be regarded as an extension of ASCII, is often the default. The current trend is to avoid giving such a special position to ISO Latin 1 among the variety of encodings.

Definitions

The following definitions are not universally accepted and used. In fact, one of the greatest causes of confusion around character set issues is that terminology varies and is often confusing.

character repertoire A set of distinct characters. No specific internal presentation in computers or data transfer is assumed. A character repertoire is usually defined by specifying names of characters and a sample (or reference) presentation of characters in visible form. Notice that a character repertoire may contain characters which *look* the same in some presentations but are regarded as logically distinct, such as Latin uppercase A, Cyrillic uppercase A, and Greek uppercase alpha.

character code A mapping, often presented in tabular form, which defines one-to-one correspondence between characters in a character repertoire and a set of nonnegative integers. That is, it assigns a unique numerical code, a code position, to each character in the repertoire. As synonyms for “code position”, the following terms are also in use: code number, code value, code element, code point, code set value. Note: The set of nonnegative integers corresponding to characters need not consist of consecutive numbers; in fact, most character codes have “holes”, such as code positions reserved for control functions or for eventual future use to be defined later.

character encoding A method (algorithm) for presenting characters in digital form by mapping sequences of code numbers of characters into sequences of octets. In the simplest case, each character is mapped to an integer in the range 0–255

according to a character code and these are used as such as octets. Naturally, this only works for character repertoires with at most 256 characters. For larger sets, more complicated encodings are needed.

glyph It is important to distinguish the character concept from the glyph concept. A glyph is a presentation of a particular shape which a character may have when rendered or displayed. For example, the character Z might be presented as a boldface **Z** or as an italic Z, and it would still be a presentation of the same character. On the other hand, lower-case z is defined to be a separate character—which in turn may have different glyph presentations.

Unicode Unicode is a standard by the Unicode Consortium (<http://www.unicode.org/>), which defines a character repertoire and character code identical with ISO 10646 and an encoding for it. In practice, people usually talk about Unicode rather than ISO 10646, partly because Unicode is more explicit about the *meanings* of characters, partly because Unicode charts of characters are available on the Web. (Unicode version 1.0 used somewhat different names for some characters than ISO 10646. In the current Unicode version 2.1 the names have been made the same as in ISO 10646.)

Notice that a character code assumes or implicitly defines a character repertoire. A character encoding could, in principle, be viewed purely as a method of mapping a sequence of integers to a sequence of octets. However, quite often an encoding is specified in terms of a character code (and the implied character repertoire). The *logical* structure is still the following:

1. A character *repertoire* specifies a collection of characters, such as “a”, “i”, and “ä”.
2. A character *code* defines numeric codes for characters in a repertoire. For example, in the ISO 10646 character code the numeric codes for “a”, “i”, “ä”, and “‰” (per mille sign) are 97, 33, 228, and 8240. (Note: Especially the per mille sign, presenting o/oo as a single character, can be incorrectly shown on display or on paper. That would be an illustration of the symptoms of the problems we are discussing.)
3. A character *encoding* defines how sequences of numeric codes are presented as (i.e., mapped to) sequences of octets. In one possible encoding for ISO 10646, the string “a!ä%,” is presented as the following sequence of octets (using two octets for each character): 0, 97, 0, 33, 0, 228, 32, 48.

The phrase character set is used in a variety of meanings. Often it denotes just a character repertoire but it may

also refer to a character code or even to a character encoding. See, for example OII’s document *Character Set Standards* (<http://guagua.echo.lu/oii/en/chars.html>), which mentions several standards, some of which define a character code while others also specify a fixed encoding.

Quite often the choice of a character repertoire, code, or encoding is presented as the choice of a *language*. A pulldown menu in a program might be labeled “Languages”, yet consist of character encoding choices (only). A language setting is quite distinct from character issues, although naturally each language has its own requirements on character repertoire.

Example: ASCII

The name ASCII, originally an abbreviation for *American Standard Code for Information Interchange*, denotes an old character repertoire, code, and encoding.

In fact, the definition of ASCII also defines a set of control codes (“control characters”) such as linefeed (LF) and escape (ESC). But the *character repertoire* proper, consisting of the *printable* characters of ASCII, is the following (where the first item is the blank, or space, character):

ASCII Printable characters

```
! " # $ % & ' ( ) * + , - . /
0 1 2 3 4 5 6 7 8 9 : ; < = > ?
@ A B C D E F G H I J K L M N O
P Q R S T U V W X Y Z [ \ ] ^ _
` a b c d e f g h i j k l m n o
p q r s t u v w x y z { | } ~
```

There are actually several national variants of ASCII. In such variants, some special characters have been replaced by national letters (and other symbols). There is great variation here, and even within one country and for one language there might be different variants. The original ASCII is therefore often referred to as US-ASCII; the formal standard (issued by ANSI, see the URL <http://www.ansi.org/>) is *ANSI X3.4-1986*.

The international standard *ISO 646* (see, for instance, <http://guagua.echo.lu/oii/en/chars.html#ISO646>) defines a character set similar to US-ASCII, but with code positions corresponding to US-ASCII characters @[\{|} as “national use positions”. It also takes some liberties with the characters # \$ ^ ‘ ~. The standard also defines an “international reference version (IRV)”, which is (in the 1991 edition of ISO 646) identical to US-ASCII.

The following table lists ASCII characters which might be replaced by other characters in national variants of ASCII. (That is, the code positions of these US-ASCII characters might be occupied by other characters needed for national use.) The lists of character appearing in national variants are not intended to be exhaustive, just typical *examples*.

dec	hex	glyph and Unicode name	variants
35	23	# number sign	# £ Û
36	24	\$ dollar sign	ⱡ
64	40	@ commercial at	É § Ä à ¢
91	5B	[left square bracket	Å Æ ° â ; ÿ é!
92	5C	\ reverse solidus	Ü Ø Ç Ñ ½ ¾
93	5D	] right square bracket	Å Ü § ê é ¿
94	5E	^ circumflex accent	Û î
95	5F	_ low line	è
96	60	` grave accent	é ä µ ô ù
123	7B	{ left curly bracket	ä æ é à °
124	7C	vertical line	ö ø ù ò ñ f
125	7D	} right curly bracket	â ü è ç ¼
126	7E	~ tilde	ü - ß " û î

Almost all of the characters used in the national variants have been incorporated into ISO Latin 1. Many systems which support ISO Latin 1 in principle may still reflect the use of national variants of ASCII in some details; for example, an ASCII character might get *printed* or *displayed* according to some national variant. Thus, even “plain ASCII text” is thereby not always portable from one system or application to another.

More information about national variants and their impact can be found at:

- *Digression on national 7-bit codes* by Alan J. Flavell
(<http://ppewww.ph.gla.ac.uk/%7eflavell/iso8859/digress.html#national>)
- The ISO 646 page by Roman Czyborra
(<http://czyborra.com/charsets/iso646.html>)
- Character tables by Koichi Yasuoka
(<http://www.kudpc.kyoto-u.ac.jp/%7eyasuoka/CJK.html>)

The *character code* defined by the ASCII standard is the following: code values are assigned to characters consecutively in the order in which the characters are listed in the “ASCII Printable characters” table above (rowwise), starting from 32 (assigned to the blank) and ending up with 126 (assigned to the tilde character ~).

The *character encoding* specified by the ASCII standard is very simple, and the most obvious one for any character code where the code numbers do not exceed 255: each code number is presented as an octet with the same value. Notice that code numbers from 0 to 31 and code numbers from 127 upwards do not correspond to any printable character. (Code numbers 0–31 and 127 are reserved for control purposes. They have standardized names, but in fact their usage varies a lot.) Octets 128–255 are not used in ASCII. (This allows programs to use the first, most significant bit of an octet as a parity bit, for example.)

Sometimes the phrase “8-bit ASCII” is used. This is a misnomer (since ASCII is by definition seven bits only)

used to refer to *various* character codes which are extensions of ASCII in the following sense: the character repertoire contains ASCII as a subset, the code numbers are in the range 0–255, and the code numbers of ASCII characters equal their ASCII codes.

Another example: ISO Latin 1 and ISO 8859-1

The ISO 8859-1 standard (which is part of the ISO 8859 family of standards) defines a *character repertoire* identified as “Latin alphabet No. 1”, commonly called “ISO Latin 1”, as well as a *character code* for it. The repertoire contains the ASCII repertoire as a subset, and the code numbers for those characters are the same as in ASCII. The standard also specifies an *encoding*, which is similar to that of ASCII: each code number is presented simply as one octet.

In addition to the ASCII characters, ISO Latin 1 contains various accented characters and other letters needed for writing languages of Western Europe, and some special characters. These characters occupy code positions 160–255, and are shown below:

```

i ç £ ⱡ § ¨ © ≡ < ¬ η ® ¯
° ± ² ³ ´ µ ¶ · ¸ ¹ º » ¼ ½ ¾ ¿
À Á Â Ã Ä Å Æ Ç È É Ê Ë Ì Í Î Ï
Ð Ñ Ò Ó Ô Õ Ö × Ø Ù Ú Û Ü Ý Þ ß
à á â ã ä å æ ç è é ê ë ì í î ï
ð ñ ò ó ô õ ö ÷ ø ù ú û ü ý þ ÿ

```

Notes:

- The first of the characters above appears as space; it is the so-called “no-break space”.
- The presentation of some characters in copies of this document may be defective, e.g., due to lack of font support. You may wish to compare the presentation of the characters on your browser with the character table presented as a GIF image (<http://czyborra.com/charsets/iso8859-1.gif>) in the famous *ISO 8859 Alphabet Soup* (<http://czyborra.com/charsets/iso8859.html>) document. (In text only mode, you may wish to use the simple table of ISO Latin 1 which contains the names of the characters.)
- Naturally, the appearance of characters varies from one font to another.

Typing characters

Typing characters on a computer may appear deceptively simple: you press a key labeled “A”, and the character “A” appears on the screen. You also expect “A” to be included into a disk file when you save what you are typing, you expect “A” to appear on paper if you print your text, and you expect “A” to be sent if you send your product by E-mail or something like that.

Thus far, you should have learned that the presentation of a character in computer storage or disk or in data transfer may vary a lot. You have probably realized that especially if it is not the common “A” but something more special (say, an “A” with an accent), strange things might happen, especially if data is not accompanied with adequate information about its encoding.

But you might still be too confident. You probably expect that on *your* system at least things are simpler than that. If you use your very own very personal computer and press the key labeled “A” on *its* keyboard, then should not it be evident that in *its* storage and processor, on *its* disk, on *its* screen its invariably “A”? Cannot you just ignore its internal character code and character encoding? Well, probably yes—with “A”. I wouldn’t be so sure about “Ä”, for instance.

When you press a key on your keyboard, then what actually happens is this. The keyboard sends the code of a character to the processor. The processor then, in addition to storing the data internally somewhere, normally sends it to the display device. Now, the *keyboard settings* and the *display settings* might be different from what you expect. Even if a key is labeled “Ä”, it might send something else than the code of “Ä” in the character code used in your computer. Similarly, the display device, upon receiving such a code, might be set to display something different. Such mismatches are usually undesirable, but they are definitely *possible*.

Moreover, there are often *keyboard restrictions*. If your computer uses internally, say, ISO Latin 1 character repertoire, you probably will not find keys for all 191 characters in it on your keyboard.

Therefore, you often need *program or operating system specific ways* of entering characters from a keyboard, either because there is no key for a character you need or there is but it does not work (properly). Two important *examples* of such ways:

- On *Windows* systems, you can produce any character in the Windows character set (naturally, in its Windows encoding) as follows: Press down the Alt key and keep it down. Then type, using the separate numeric keypad (not the numbers above the letter keys!), the four-digit code of the character in decimal. Finally release the Alt key. Notice that the first digit is always 0, since the code values are in the range 32–255 (decimal). For instance, to produce the letter “Ä” (which has code 196 in decimal), you would press Alt down, type 0196 and then release Alt. Upon releasing Alt, the character should appear on the screen. This method is often referred to as **Alt-*nnnn***. (Notice that if you omit the leading zero (i.e., use **Alt-*nnn***), the effect is *different*, since that way you insert the character in code position *nnn* in octal in the *DOS character code*! Quite confusing, since the presence or absence of the leading zero affects *both the number base*—octal or decimal—for interpreting the digit

sequence as a number *and the encoding* used to map the number into a character. For example, **Alt-196** would probably insert a graphic which looks like a hyphen.)

- In the *Emacs* editor in Unix, you can produce any ISO Latin 1 character by typing first **control-Q**, then its code as a three-digit octal number. To produce “Ä”, you would thus type **control-Q** followed by the three digits 304 (and expect the “Ä” character to appear on screen). This method is often referred to as “**~Q-*nnn***”. There exist other ways of entering many ISO Latin 1 characters in Emacs, see the URL

<http://www.hut.fi/u/jkorpela/emacs-iso.html>.

It is often possible to use various “escape” notations for characters. This somewhat vague term means notations which are converted to characters according to some specific rules by some programs. For example, in the HTML language one can use the notation **Ä** or, equivalently, the notation **Ä** for the character Ä. To take another example, in the C programming language, one can usually write **\0304** to denote “Ä” within a string constant, although this makes the program character code dependent. In cases like these, the character itself does not occur in a file (such as an HTML document or a C source program); instead, the file contains the “escape” notation as a character sequence, which will then be *interpreted* in a specific way by programs like a Web browser or a C compiler.

Information about encoding

The need for information about encoding

It is hopefully obvious from the preceding discussion that *a sequence of octets can be interpreted in a multitude of ways* when processed as character data. By looking at the octet sequence only, you cannot even know whether each octet presents one character or just part of a two-octet presentation of a character, or something more complicated.

Naturally, such a sequence could be intended to present other than character data, too; it could be an image in a bitmap format, or a computer program in binary form, or numeric data in the internal format used in computers.

This problem can be handled in different ways in different systems when data is stored and processed within one computer system. For *data transmission*, a platform-independent method of specifying the encoding and other relevant information is needed. Such methods exist, although they not always used widely enough. People still send each other data without specifying the encoding, and this may cause a lot of harm. Attaching a human-readable note, such as a few words of explanation in an E-mail message body, is better than nothing. But since data is processed by programs which

cannot understand such notes, the encoding should be specified in a standardized computer-readable form.

The MIME solution

Internet media types, often called MIME media types, can be used to specify a major media type (“top level media type”, such as `text`), a subtype (such as `html`), and an encoding (such as `iso-8859-1`). They were originally developed to allow sending other than plain ASCII data by E-mail. They can be (and should be) used for specifying the encoding when data is sent over a network, e.g., by E-mail or using the HTTP protocol (see <http://www.w3.org/pub/WWW/Protocols/>) on the World Wide Web.

The media type concept is defined in RFC 2046 (<ftp://nis.nsf.net/internet/documents/rfc/rfc2046.txt>). The procedure for registering types is given in RFC 2048 (<ftp://nis.nsf.net/internet/documents/rfc/rfc2048.txt>). For less authoritative but more readably presented information on media types, see, e.g., Chris Herboth’s *MIME Types* Web page at the URL <http://www.qnx.com/%7Echrish/Be/info/mime-types.html>. The technical term used to denote a character encoding in the Internet media type context is *character set*, abbreviated *charset*.

Specifically, when data is sent in MIME format, the media type and encoding are specified in a manner illustrated by the following example:

```
Content-Type: text/html; charset=iso-8859-1
```

This specifies, in addition to saying that the media type is `text` and subtype is `html`, that the character encoding is ISO 8859-1.

Several character encodings have alternate names in the registry. For example, the basic (ISO 646) variant of ASCII can be called “ASCII” or “ANSI_X3.4-1968” or “cp367” (plus a few other names); the preferred name in MIME context is, according to the registry, “US-ASCII”. Similarly, ISO 8859-1 has several names, the preferred MIME name being “ISO-8859-1”. The “native” encoding for Unicode, UCS-2, is named “ISO-10646-UCS-2” there.

The `Content-Type` information is an example of information in a header. Headers relate to some data, describing its presentation and other things, but is passed as logically separate from it. Possible headers and their contents are defined in the basic MIME specification, RFC 2045 (<ftp://nis.nsf.net/internet/documents/rfc/rfc2045.txt>). Adequate headers should normally be generated automatically by the software which sends the data (such as a program for sending E-mail, or a Web server) and interpreted automatically by receiving software (such as a program for reading E-mail, or a Web browser). In E-mail messages, headers precede the message body; it depends on the E-mail program whether and how it displays the headers. For

Web documents, a Web server is required to send headers when it delivers a document to a browser (or other user agent) which has sent a request for the document.

How MIME should work in practice

Basically, MIME should let people communicate smoothly without hindrances caused by character code and encoding differences. MIME should handle the necessary conversions automatically and invisibly.

For example, when person A sends E-mail to person B, the following should happen: The E-mail program used by A encodes A’s message in some particular manner, probably according to some convention which is normal on the system where the program is used (such as ISO 8859-1) encoding on a typical modern Unix system). The program automatically includes information about this encoding into an E-mail header, which is usually invisible both when sending and when reading the message. The message, with the headers, is then delivered, through network connections, to B’s system. When B uses his E-mail program (which may be very different from A’s) to read the message, the program should automatically pick up the information about the encoding as specified in a header and interpret the message body according to it. For example, if B is using a Macintosh computer, the program would automatically convert the message into Mac’s internal character encoding and only then display it. Thus, if the message was ISO 8859-1 encoded and contained the “Ä” (upper case A with dieresis) character, encoded as octet 196, the E-mail program used on the Mac should use a conversion table to map this to octet 128, which is the encoding for “Ä” on Mac. (If the program fails to do such a conversion, strange things will happen. ASCII characters would be displayed correctly, since they have the same codes in both encodings, but instead of “Ä”, the character corresponding to octet 196 in Mac encoding would appear - a symbol which looks like f in italics.)

Problems with implementations—examples

Unfortunately, there are deficiencies and errors in software so that *users* often have to struggle with character code conversion problems, perhaps correcting the actions taken by programs.

Typical minor (!) problems which may occur in communication in Western European languages other than English is that most characters get interpreted and displayed correctly but some “national letters” are rendered wrongly. For example, the character repertoire needed in German, Swedish, and Finnish is essentially ASCII plus a few letters like “ä” from the rest of ISO Latin 1. If a text in such a language is processed so that a necessary conversion is not applied, or an incorrect conversion is applied, the result might be that, e.g., the word “später” becomes “spter”, or “sp1ter”, or “spdter”,

or even “sp=E4ter”.

To illustrate what may happen when text is sent in a grossly invalid form, consider the following example. I am sending myself E-mail, using Netscape 4.0 (on Windows 95). In the mail composition window, I set the encoding to UTF-8. The body of my message is simply

Tämä on testi.

(That’s Finnish for “This is a test”. The second and fourth character is the letter a with an umlaut.) Trying to read the mail on my Unix account, using the Pine E-mail program (popular among Unix users), I see the following (when in “full headers” mode; irrelevant headers omitted here):

```
X-Mailer: Mozilla 4.0 [en] (Win95; I)
MIME-Version: 1.0
To: Jukka.Korpela@hut.fi
Subject: Test
X-Priority: 3 (Normal)
Content-Type: text/plain; charset=x-UNICODE-2-0-UTF-7
Content-Transfer-Encoding: 7bit
```

```
[The following text is in the
      "x-UNICODE-2-0-UTF-7" character set]
[Your display is set for the
      "ISO-8859-1" character set]
[Some characters may be displayed incorrectly]
```

T+06Q- on testi.

Interesting, is it not? I specifically requested UTF-8 encoding, but Netscape used UTF-7. And it did not include a correct header, since x-UNICODE-2-0-UTF-7 is not a registered “charset” name. Even if the encoding had been a registered one, there would have been no guarantee that my E-mail program would have been able to handle the encoding. The example, “T+06Q-” instead of “Tämä”, illustrates what may happen when an octet sequence is interpreted according to another encoding than the intended one. In fact, it is difficult to say what Netscape was really doing, since it seems to encode incorrectly.

A *correct* UTF-7 encoding for “Tämä” would be “T+AQQ-m+AQQ-”. The “+” and “-” characters correspond to octets indicating a switch to “shifted encoding” and back from it. The shifted encoding is based on presenting Unicode values first as 16-bit binary integers, then regrouping the bits and presenting the resulting six-bit groups as octets according to a table specified in RFC 2045 (<ftp://nis.nsf.net/internet/documents/rfc/rfc2045.txt>) in the section on Base64. See also RFC 2152 (<ftp://nis.nsf.net/internet/documents/rfc/rfc2152.txt>).

Practical conclusions

Whenever text data is sent over a network, the sender and the recipient should have a joint agreement on the

character encoding used. In the optimal case, this is handled by the software automatically, but in reality the users need to take some precautions.

Most importantly, make sure that any Internet software that you use to send data specifies the encoding correctly in suitable headers. There are two things involved: the header must be there and it must reflect the actual encoding used; and the encoding used must be one that is widely understood by the (potential) recipients’ software. (One must often make compromises as regards to latter aim: you may need to use an encoding which is not yet widely supported to get your message through at all.

Learn to use your Web browser, newsreader, and E-mail program so that you can retrieve the encoding information for the page, article, or message you are reading. (For example, on Netscape use **View Page Info**; on News Xpress, use **View Raw Format**; on Pine, use **h**).

If you use, say, Netscape to send E-mail or to post to Usenet news, make sure it sends the message in a reasonable form. In particular, make sure it does not duplicate the message by sending it both as plain text and as HTML (turn off the latter). As regards to character encoding, make sure it is something widely understood, such as ASCII, some ISO 8859 encoding, or UTF-8, depending on how large a character repertoire you need.

In particular, avoid sending data in a proprietary encoding (like the Macintosh encoding or a DOS encoding) to a public network. (At the very least, if you do that, make sure that the message heading specifies the encoding!) There is nothing wrong with using such an encoding within a single computer or in data transfer between similar computers. But when sent to Internet, data should be converted to a more widely known encoding, by the sending program. If you cannot find a way to configure your program to do that, get another program.

Further reading

- *Character Sets*
(<http://www.ewos.be/tg-cs/gtop.htm>) in *GOSS*
(<http://www.ewos.be/goss/top.htm>).
- *Character Encodings Concepts*
(<http://developer.apple.com/techpubs/mac/TextEncodingCMgr/TECRefBook-140.html>), adapted from a presentation by Peter Edberg at a Unicode conference. A rich source of information, with good illustrations.
- Alan J. Flavell’s *ISO-8859 briefing and resources*
(<http://ppewww.ph.gla.ac.uk/%7eflavell/iso8859/iso8859-pointers.html>). Partly a character set tutorial, partly a discussion of specific (especially ISO 8859 and HTML related) issues in depth.

- *Guide to the use of character sets in Europe* (<http://www.stri.is/TC304/GUIDE/>). A draft which contains explanations of basic concepts related to character sets in general and discusses various European standards.
- *Character Set Standards* (<http://guagua.echo.lu/oii/en/chars.html>). A lot of technical data can be found at <http://guagua.echo.lu/oii/en/oii-home.html>.
- Usenet FAQs related to internationalization (relatively out of date) (<http://www.cis.ohio-state.edu/hypertext/faq/usenet/internationalization/top.html>).
- “Character Set” Considered Harmful (<http://www.w3.org/MarkUp/html-spec/charset-harmful.html>) by Dan Connolly (<http://www.w3.org/People/Connolly/>). A good discussion of basic concepts and misconceptions.
- *The Nature of Linguistic Data: Multilingual Computing* (<http://linguist.emich.edu/%7elinglit/routledge/multilingual.html>) a collection of annotated links to information on character codes, fonts, etc.
- Michael Everson’s Web site (<http://www.indigo.ie/egt/>) It contains a lot of links to detailed documents on character code issues, especially progress and proposals in standardization.
- An online character database by Indrek Hein (<http://www.eki.ee/letter/>) at the Institute of the Estonian Language (<http://www.eki.ee/>). You can for instance search for Unicode characters by name or code position, get lists of differences between some character sets, and get lists of characters needed for different languages.

6.2 Questions and Answers from the Computing Help Desk

Collected by the IT/User Support Team

This is a collection of Questions & Answers that have been treated by the Computing Help Desk managed by the User Support group.

Some PC Tips

Hannes Schwarzbauer, IT/User Support

Question – Microsoft “Windows” symbol

I have a modern PC keyboard which has a key with the Microsoft “Windows” symbol on it. It is beneath the space and alt key. What is it good for?

Answer

It is a “program function” or “shortcut” key, which lets you rapidly execute some program or call up some functions.

Note, however, that this key is **not** available on all standard keyboards.

Here is a list of available functions:

Win Key (WK):

=====

WK + E	Open/Ouvrir [Explorer]
WK + R	Open/Ouvrir [command Run]
WK + Pause	Open/Ouvrir [System Properties]
WK + F	Display/Afficher [Find: All Files]
WK + Ctrl + F	Display/Afficher [Find : Computer]
WK + M	Minimize/Minimiser active windows
WK + Shift + M	Discard/Annuler active windows
WK + Tab	Go through/Circuler dans [task bar]
WK + F1 ou F1	Display/Afficher [Help Windows]

Question – Error messages from “Doctor Watson”

When I launch various PC applications, I keep getting error messages from Doctor Watson. What does this mean? Who is this doctor?

Answer

Dr. Watson is a problem report generator that kicks in when an application does something really bad. In general it means troubles, like memory problems.

If the user is running a standard NICE PC, the only reason for this to happen is a hardware problem. It can be anything from a faulty memory card to a dead processor fan with the heat rising on the motherboard.

If “rechecking everything” fails, the next option is to reinstall the system. If even this fails, it definitely is a hardware problem and the user must contact his normal PC support team, or the original supplier (for example the IT PC Desktop Service if the PC was originally purchased there).

Question – passwords

I have changed my password but it only works on W95, not on NT. Are they not synchronised? The error message I get is: “Your password change request HAS NOT been processed because of an internal error. Please try again later”.

Answer

This happens when a user's password is too long. It can be considered as a "bug" in NT. NT only allows a maximum of 15 characters. The user must not use more than this maximum.

Question – Application log file full

Under Windows NT, I get the message: "Application log file full". What can I do?

Answer

This log is part of the NT event viewer. To clear it, go from "programs" into "administrative tools", then "event viewer". Before you clear the log and to be on the safe side, save the old one as whatever you want (for instance, oldfile).

Note that you have to be logged as Administrator to do the above

Question – Automatically Adjust Clock

If you are on a network which operating system has the functionality of adjusting for changes in daylight savings time, there will be a problem if you also use the Automatically Adjust Clock for Daylight Savings Time Changes feature in Windows 95, particularly if all workstations on the network are either not using Windows 95 or have disabled this feature in Control Panel, Date/Time, Time Zone tab.

If you allow the network operating system as well as Windows 95 to adjust for these changes, then you are no longer synchronized with the rest of the network and entries made into other users' schedules will be changed to try to adjust for this.

The following is an example of a Schedule+ problem that has been encountered: a Windows 95 user used Schedule+ to schedule meetings on a Novell network, running Windows 95 and Office 95. The other workstations are running Windows for Workgroups. The network operating system had adjusted all clocks on the system when there was a change between daylight saving time and eastern standard time. When the Windows 95 user was asked if he or she wanted to have Windows 95 automatically adjust for changing the clocks, the user chose "Yes". After this, all other schedules on the network that had meetings scheduled, prior to the time change, for 1:00 P.M. found that the time has now been changed to 2:00 P.M.

Explanation: if the network operating system adjusts the time and then Windows 95 adjusts the time, the Windows 95 machine is no longer synchronized with the other machines on the network. It has actually had its time set back by one hour twice. The workstation appears to be in a different time zone than the rest of the machines. In an effort to synchronize the times, the scheduled appointment times change.

Answer

To avoid this problem, choose Date/Time from the Control Panel. In the Time Zone tab, disable Automatically adjust clock for daylight savings time changes.

More in particular, if you click the question mark (?) on Automatically adjust clock for daylight savings time changes, you get an explanation of the feature as well as the following warning: "If you use more than one operating system, make sure that only one of them adjusts for daylight saving time"

Also refer to the question on AFS clock skew on UNIX below

Registry Settings—important remark

The two following questions ("Deactivate screen saver" and "'pagefile.sys' file") imply manipulations of the registry; please note that users must understand what they do before changing any value! Otherwise undesired side effects—that maybe difficult to debug even by the NICE experts—may occur afterwards.

Question – Deactivate screen saver

When a password protected screen saver starts, is it possible to get some grace period before a password needs to be entered to deactivate?

Answer

Yes this is possible. Although not set by default in NICE, this can be changed as follows:

1. Start the registry editor (regedit.exe).
2. Move to
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows
(or Windows NT) \CurrentVersion\Winlogon
3. From the edit menu select New - String Value.
4. Enter a name of ScreenSaverGracePeriod and press Enter.
5. Double click and set to the number of seconds grace period. Click ok.
6. Close the registry editor.
7. Reboot the machine.

Question – The 'pagefile.sys' file

I have recently ran short on disk space on my hard drive on NT. By inspecting files I discovered a huge (in my case over 40MB) file called "pagefile.sys". Is it possible to safely delete or at least shrink this file?

Answer

During an NT session, it is not possible to access or open this system file. However, it can be shrunk/cleaned by the following procedure:

1. launch the registry editor (regedt32.exe or regedit).
2. Change to one (1) the value of the key
ClearPageFileAtShutdown, which is in:
HKEY_LOCAL_MACHINE\System\CurrentControlSet
\Control\SessionManager\Memory Management

If this key does not exist, you have to create it:

```
Value Name: ClearPageFileAtShutdown
Value type: REG_DWORD
Value: 1
```

You then have to reboot your NT machine to make then changes take effect.

Note that on an extremely full hard drive, it may be necessary to reduce the virtual memory size to a value less than the drive space available:

- Control panel - system
- performance - virtual memory

Some Unix Tips

Nicole Crémel and Roger Woolnough, IT/User Support

Question – AFS Cache

When listing the directory content on 2 different nodes of the ATLAS server I got 2 different versions (one missing the latest corrections):

```
atlas03 > rsh atlas02 ll
/afs/cern.ch/atlas/users/pop/.../MuonIdNtuple.cxx
-rw-r--r--  1 pop    zp   13390 Mar 26 14:06
atlas03 > rsh atlas03 ll
/afs/cern.ch/atlas/users/pop/.../MuonIdNtuple.cxx
-rw-r--r--  1 pop    zp   13197 Mar 21 19:41
```

What can be done?

Answer

You can rectify this abnormal situation with the command `fs flush` which forces the AFS Cache Manager to remove each specified directory or file from its cache of data and status information. The result is that the next time data from a flushed directory or file is requested, the Cache Manager contacts the File Server for the most current version, along with a new callback (if necessary) and associated status information. This command does not discard data from application program buffers or data that has been altered in the cache but not yet written back to the central copy maintained by the File Server.

You can get all details regarding this AFS command with `man fs_flush`.

Question – Work Group Server

Where to get all information (responsible, user guide, configuration details, performance reports, etc.) on a specific Work Group Server? (e.g., THWGS, CMS, Atlas, Alice, Chorus, NA48, NA49, TIS).

Answer

At the URL <http://wwwinfo.cern.ch/pdp/bis/main/experiments.html> you will find links, the name of the “Local” and “IT” contact person, and the mailing list corresponding to the Work Group Server for all experiments.

Question – fs sysname

Why do on Rsplus (Aix 4.3) the 2 commands `fs sysname` and `uname -r` (or `uname -a`) provide different information:

```
[rsplus11] ~ > fs sysname
Current sysname is 'rs_aix42'
[rsplus11] ~ > uname -r
3
[rsplus11] ~ > uname -a
AIX rsplus11 3 4 001520FD4C00
```

Answer

The AFS command `sysname` command is a relatively ad hoc constructions used basically to discriminate between different AFS binaries. It is used because it is automatically translated from `@sys` when used in an AFS path. However the exact “system release” is effectively given by the UNIX command `uname -r`.

Question – StarOffice (Word, Excel and Powerpoint)

I would like to use the utility StarOffice, is it available at CERN through AFS?

Answer

We know that many users are interested by this product which contains *Word*, *Excel* and *Powerpoint 97* look-alikes. CERN is currently negotiating to have a site-wide licence and expects to have something in place shortly.

Note that the first versions will be for Solaris and Linux. Both *Word* and *Excel* utilities are fairly well compatible with the PC/Mac versions but the *Powerpoint* version is not able to re-transfer files from Linux/Solaris back to PC or Mac, although it can read them. In the future there may be versions available for HP, but not before the end 1999, at the earliest.

Question – “AFS clock skew” problem

I installed the Transarc AFS client for NT and wanted to get an AFS token afterwards with `afs_auth.exe` but I got the following error message: “server and client clocks are badly skewed”.

I use GMT+1:00 for the time zone. Do you know how to solve this problem?

Answer

The “AFS clock skew” problem has been seen elsewhere on several occasions, fortunately in all cases it has come down to the same quirk in the Windows/NT setup:

1. go to the “clock settings” panel (double-click the clock);
2. set your time zone to GMT , hit “Apply” and;
3. set the time to “now” minus 1 hour;
4. set your time zone back as you wish;
5. hit “Apply” and verify that now the time is correct;
6. reboot (this may be not necessary, but preferable);
7. try the “AFS authentication” again.

Some DeskTop Publishing Tips

Andreas Wagner, IT/User Support

Question – Printing national characters

I would like to print a text file containing national characters ? When using `xprint` I only get strange symbols, what should I do?

Answer

You can use the program `a2ps` to print such files, for instance, typing `a2ps -P 513-us filename`.

`xprint` converts text files to PostScript by using internally an older version of `a2ps`. This version does not handle national characters properly. It is planned to update `xprint` in the future.

Miscellaneous

Question – Escape key on and NCD terminal

I am using a NCD terminal, and there is at least one key (escape) that does not work properly. I do not mean that the key is mechanically bad, it just does not do what is supposed to do. How do I fix it?

Answer

If you have no pure “Esc” key and have an “F11” key labeled also “(Esc)” you can assign the value of “Esc” to this key by issuing the command: `xmodmap -e "keysym F11 = Escape"`

You can then enter this command into some startup rc-file.

Question – Ethernet address

What is the “Ethernet address” of a machine and how to find it?

Answer

The ethernet address is 12 hexadecimal (0–9, A–F) digits, of which the first six specify the maker of the interface and the last six are unique. On some systems this address is already printed on a label, either outside the system or inside on the interface card which has the network connector.

On a PC running Windows 95 there is normally a Control Panel program (menu **Start - Settings -**

Control Panel) called “TCP/IP”. Executing this shows a window giving the IP address and name, plus a small icon labelled “View Details”. Clicking on this small icon should then give a second window including a field called Adapter Address, which specifies the ethernet Adapter Address.

For a PC running NT it is necessary first to open a Command Prompt window. On a NICE NT PC this is in “PROGRAMS” in the START menu. The command `wntipcfg` will then normally give the same window as for the Windows 95 PC, including the ethernet Adapter Address. If this `wntipcfg` command is unknown (perhaps because the NT is stand-alone) an alternative command is `net config work`, which will produce a line like:

Workstation active on NetBT_1EEPR01 (Adapter Address)

On Unix systems the address is indicated at boot time. It may also be possible to issue one of the following commands:

<code>netstat -i</code>	(on IBM AIX and DEC : needs the interface running!)
<code>/usr/sbin/lscfg -v</code>	(on AIX)
<code>dmesg</code>	(on Linux and Solaris : normally in <code>/usr/sbin</code>)
<code>lanscan</code>	(on HP/UX : normally in <code>/usr/sbin</code>)

For a PC running Linux the ethernet address is written at startup time into a log file: normally `/var/log/syslog` or `/var/log/messages`.

On a recent Macintosh there is a Control Panel program called TCPIP. Executing this shows a window giving the IP address. When this is visible there is an item in the “File” menu called “Show Info”. Using this will normally give a separate window showing the ethernet address. If it does not then you have an old version of what is now called “Open Transport”. To update choose the File Server `Srv0_Nice` in zone `NOVELL`, select log on method “Apple Standard UAMs”, enter as Guest, go into the Folder “Applications”, then the Folder “System Tools”, then the Folder “Universal System”, then the Folder “OT 1.1.2-NET Install” and then double click on “Open Transport Installation”.

Older Macintoshes have in the Control Panel a program called MacTCP. Executing this should show a window including an icon called “Ethernet”. Holding down the “Option” key while clicking on this icon should then show the ethernet address. Note that this “Option” is sometimes marked “alt” and may have a little diagram on it, such as:

```

----
 \__

```

The network team at CERN needs to know the exact ethernet address for a number of reasons. Primarily, it enables them to detect if someone else tries to “borrow” and use the corresponding IP address. In addition, it may happen that they detect hardware or software problems related to the ethernet address and wish to inform the owner and/or responsible person. The information is particularly valuable for systems registered as portable and not necessarily having a fixed home network socket or IP name and address.

If you are working at CERN you should be able to check all the networking parameters of your device (and eventually inform the network team of any wrong or missing information) by using the web form at the URL <http://network.cern.ch/register>

Question – LHC++

What are the definition for LHC++ software both at CERN and from outside, and especially the AFS availability?

Answer

Everything is explained in the LHC++ Web page at the URL <http://wwwinfo.cern.ch/asd/lhc++/access.html>

Regarding access from AFS (the root directory is at `/afs/cern.ch/sw/lhcxx`) anyone from a cern node can have general access. From an outside node, only the person who was designated to be in the `cern:li` group when registering can have access to the software. And in this case, only to the export tree: `/afs/cern.ch/sw/lhcxx/export` for exporting the software off-site. The policy is not to allow direct use from afs for off-site. To see who is in the `cern:li` group type `pts mem cern:li`.

Question – The Java programming language

Who can I contact at CERN for questions/problems with the Java programming language?

Answer

There is no official support for Java but we invite users to address their questions/problems to the `cern.java` newsgroup.

This newsgroup is intended to be the most basic communication channel for the community of Java developers at CERN. In order not to restrict the immediate evolution the scope is initially quite wide:

- To report problems and issues regarding the current deployment of the JDK on all platforms. This includes also communications about upgrades, fixes, improvements, etc.
- To expose, agree and follow up issues about the evolution of the Java facilities at CERN or simply about Java within the CERN context.
- To hold generic discussions about Java: self assistance by the developer community, issues about the language itself, etc.

We expect to structure the newsgroup (`cern.java.*`) in a consensual way as needs arise and different elements of the Java infrastructure settle down. Also, no alternative communication mechanism is discarded (creation of a Java Users group, etc..) as long as it is convenient for the community and there is consensus to establish it.

Question – Web access with AFS files

I would like to restrict read access to certain web pages, and have protected them as far as the access via the web is concerned. The files themselves are on AFS.

How can I make them accessible to the server, but protect them from being read by other AFS users?

Answer

For the three Web servers `consult.cern.ch`, `home.cern.ch` and `wwwinfo.cern.ch` this can be achieved using AFS group names in a “generic” way.

- For `consult` give read access to `webserver:consult`
- For `home` give read access to `webserver:home`
- For `wwwinfo` give read access to `webserver:wwwinfo`

For example, if you want to restrict AFS read access to one of your home directories, just type:

```
fs setacl mydirname webserver:home read
```

7. User Documentation

Everything related to “computer documentation”: updates to computer documentation distributed by IT (including the UCO references cards), book reviews, offerings of the COBS service, etc.

7.1 News from the User Support Bookshop

Roger Woolnough and Jutta Megies, IT/User Support

We should start by informing those persons who are not always in the Computer Centre that we have installed a display case next to the User's Office in Building 61 next to the public terminals. We also offer a 'browsing' service in Building 40 where a new satellite helpdesk has started (see the article on page 1), where the most popular titles and our latest editions can be seen. In the future this service could be expanded if we find it a success. Please send us your comments to cobs.service@cern.ch.

We have received the following books since the last CNL:

- *Classroom in a Book, Photoshop 5.0*: This book is in the series of tutorials which we have for *Framemaker* and is very popular. A similar title was taken for *Photoshop* by popular demand.
- *Windows NT Workstation 4.0 User Manual*: This is a QUE series book like the *Using* series and has a very practical approach for new and intermediate users of Windows NT. Again this book has been requested by users starting NT and requiring instructions and documentation.
- *Advanced Corba Programming with C++*: A specialised title from the *Addison Wesley Professional Series*.
- *The Mathematica Book* by Stephen Wolfram: Although Version 4 is only just released in the US this new edition of the famous book covers all the changes. Now in hard-back (and as big as ever!) this book covers all you need to know on *Mathematica*.
- *Word/Excel 97 Clics et Declics* (en français): Les deux livres *Comment faire* n'étant plus disponibles, ces deux petits livres vous expliquent l'essentiel afin d'être immédiatement opérationnel sur *Word* et *Excel*.
- *Linux Unleashed*: This SAMS publishing book is now updated to version 5 of the *Red Hat* software

and has been requested as the translation of the French Linux book that we have in stock. A very comprehensive guide covering all aspects of Linux.

- *Linux Resources d'Experts* (en français): traduction du livre précédent (*Linux Unleashed*), pour répondre à la demande de livres sur Linux en français. Basé sur les normes de *Red Hat* ce livre est très complet, pour le débutant comme pour l'expert.
- *Filemaker Pro 4, a Developer's Guide*: a long awaited serious book on this utility which, although popular, has had few books published on it. This book covers all aspects and is produced as a tutorial. Includes a CD with examples. The book also covers publishing databases on the WWW.

We have new versions of the books *Designing Web Graphics 3* and also the *Postscript Reference Guide*. We have a new version of the Java Tutorial with a follow up title *Java Tutorial Continued*.

In the coming weeks we will receive a new title from our IT *T_EX/L^AT_EX/HTML/XML* specialist Michel Goossens, entitled *The Latex Web Companion, Integrating TeX, HTML and XML*. Again by popular demand we have *The Dreamweaver Bible* which covers this alternative web authoring tool: this should be in stock by the time the CNL is published.

We are informed by *Pearson Educational* (this is the new name for the Group that includes *Addison Wesley Longman*, *Prentice Hall*, *Simon Schuster*, *Financial Times Management* etc.) that *Simulink* is no longer available as a package. We have four copies of this software left that sells with *Matlab* and once they are sold we can get no more.

Please do not hesitate to contact us with any suggestions for catalogue replacements or additions.

7.2 The UCO Reference Cards - Current List

Roger Woolnough, IT/User Support

The list is compiled regularly from the "Writeups at CERN" Web page, at URL:

<http://consult.cern.ch/writeups/S/RefCatList.html>

Computer Registration Documents

- 01 CERN Computer Centre User Registration Form
- 02 CERN Computer Centre User Registration Form
- 03 CERN Computer Center User Registration (Super Computer)
- 04 Request for Account Modification/Deletion
- 05 Règlement révisé pour l'utilisation des services informatiques au CERN
- 06 Revised rules for use of computing facilities at CERN

Electronic Mail

- 23 Intro to Pine

Text Processing

- 20 Include ps and eps files into L^AT_EX

Text Editors

- 15 vi editor
- 16 pico editor
- 17 Emacs mini reference
- 34 GNU Emacs reference card

Programming Languages

- 22 SQL Basic introduction

Unix

- 09 Should I have a dot (".") in my PATH

- 10 Customising your shell – zsh specific
- 11 Customising your shell – tcsh specific
- 12 Customising your shell – zsh options
- 14 Unix Introduction
- 41 UNIX News Reader Survey

Browsers

- 18 xdvi and ghostview
- 28 Lynx browser

Software Packages

- 26 EDH (French)
- 27 EDH (English)

VAX/VMS

- 31 VAX/VMS Card

Miscellaneous

- 999 AFS Quick Reference
- 999 LSF Quick Reference (Load Sharing Facility)
- 13 UNIX connectivity from Falco terms
- 19 make command
- 36 Zephyr Messaging Service for Unix Users
- 39 Importing CERN Computer Equipment into France
- 40 Obtaining and Configuring a Modem to Connect to CERN from Home
- 51 Using Tape Services
- 52 Zephyr Messaging for Windows Users

7.3 User Support Book Catalogue

Michel Franceschi and Roger Woolnough, IT/User Support

Situated in Building 513 1-022 (tel. 74050) is the User Support Bookshop 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 cobs.service@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 the URL:

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