


CREMEL
Nicole Laurence

RES-S/IT-US

CERN COMPUTER NEWSLETTER



CERN - IT Division
CH-1211 Geneve 23
SWITZERLAND

JANUARY – MARCH 1998

No. 230

Contents

If you need Help	i
Letters from/to the Editor	1
1 General	2
1.1 The 1998 CERN School of Computing	2
1.2 The RSPLUS Public Login Service	3
1.3 The Public Batch Service: An Upgrade and the Migration to LSF	4
1.4 Migration of 3480s to Redwood: Progress Satisfactory	5
1.5 The Future of VXCERN	5
1.6 WebOffice Tools	6
1.7 New User Registration Form	6
2 Mail and News Issues	7
2.1 Netscape 4 Mail	7
2.2 Unsolicited Electronic Mail	7
2.3 Macintosh Mail Migration	8
2.4 Pine Now Available on VXCERN	9
3 Desktop Computing	10
3.1 Changements des Droits d'Accès aux Répertoires <i>Home</i> de NICE	10
3.2 Changes of the Default Access Rights to NICE Home Directories	10
3.3 Operating System Certification - Present List	12
3.4 Operating Systems Availability on ASIS and ASIS News	13
3.5 New ASIS Mailing Lists	13
3.6 Zephyr: The Single Notification System for All Users of Central Computing Services	14
3.7 Software Development Tools Service	14
4 Communications and Networks	15
4.1 Computer Networking at CERN	15
4.2 DECnet Terminates to INFN and UK	17
5 Application Software & Databases	18
5.1 CERNLIB 98 Release	18
5.2 PAW Release Highlights	18
5.3 Replacement of the CS2-PIAF Service	18
6 Text Processing	19
6.1 New versions of $\text{\TeX}$ and $\text{\LaTeX}$ 2html	19
6.2 The Web Prepares for the Future	19
7 Questions and Answers from the UCO	24
8 Computer Documentation	27
8.1 Book News from the UCO	27
8.2 The UCO Reference Cards—Current List	27
8.3 The COBS Review - Books on <i>Java</i>	28
8.4 UCO Book Catalogue	30

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
231	Monday 1st June 1998.

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/cnl/230>

This document was produced from the HTML master and translated into L^AT_EX with the gf program. It was then somewhat edited by hand and typeset using the cerncn1 L^AT_EX class. Please note 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/230
ftp> binary
ftp> get cnl230.ps.gz (or get cnl230.ps)
ftp> quit
```

Please note that, if you do not have the gnu gunzip utility on your system you can get the uncompressed PostScript version by typing the command get cnl230.ps, without the gz suffix. In order to save Internet bandwidth, you are, however, strongly urged to try and install the gunzip utility since gzipped files are about three times smaller than their unzipped equivalents.

If you want to subscribe to the `cern-cnl-announcement` mailing list, in order to receive a mail whenever a new issue of the CERN Computer NewsLetter is available on the Web, you can fill in the form at URL:

<http://consult.cern.ch/cnls/cnl-announcement.html>

or directly send a mail (with your exact e-mail address) to the editor (Cnl.Editor@cern.ch).

If you want to receive a copy of the paper version of the CNL (at CERN or at your home institute), fill in the form at URL:

<http://consult.cern.ch/cnls/register.html>

or send an e-mail (with your exact postal address) to Cnl.Registration@cern.ch

N.B. You can cancel your subscription in both cases by filling in the appropriate form accessible from our Web page at URL:

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

Responsible editor: Nicole Cremel

Technical co-editor: Michel Goossens

IF YOU NEED HELP (Contacts at CERN)

Please use the generic electronic mail address for the service 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.

Service	Name	Address	Telephone	E-mail Address
USER SUPPORT				
User Consultancy Office (UCO)		513/R-052	74952	User.Support@cern.ch
Opening hours: Monday 10:00-12:30 and 14:00-17:30, Tuesday-Friday 9.00-12:30 and 14:00-17.30				
User Relations and Service Manager	M.Marquina	513/1-011	74912	User.Support@cern.ch
Help Desk Manager	R.Woolnough	513/1-020	79156	User.Support@cern.ch
User Registration	M.C.Perler	513/1-017	75004	User.Registration@cern.ch
Accounting Service	A.Koppanyi	513/1-019	74933	User.Support@cern.ch
Computer Book Selling Service	M.Franceschi	513/1-022	72371	COBS.Service@cern.ch
CENTRAL COMPUTER OPERATIONS				
Central Computer Operators (24/24 hours on call)		513/R-049	75011	Computer.Operations@cern.ch
CERNSP (IBM SP2/AIX)	H.Renshall	31/2-007	72377	Harry.Renshall@cern.ch
Work Group Servers	T.Cass	31/1-017	78724	Tony.Cass@cern.ch
<i>Public Login Servers (interactive UNIX servers)</i>				
- HPPLUS (HP/HP-UX)	T.Smith	31/1-013	73713	Tim.Smith@cern.ch
- DXPLUS (DEC-ALPHA/Digital UNIX)	A.Miotto	31/1-021	79576	Alessandro.Miotto@cern.ch
CORE Services				
- CSF (Simulation Cluster)	V.Doré	31/2-012	72397,13+3830	Csf.Support@cern.ch
- SHIFT (Reconstruction Cluster)	F.Hemmer	31/3-006	76104,160128	Frederic.Hemmer@cern.ch
PaRC Service (Engineering Cluster)	P.Martucci	31/2-012	78611,160277	parcsp@mail.cern.ch
PIAF Service (Interactive Analysis)	B.Ferrero	513/2-006	74886	piaf.support@cern.ch
CS-2	E.Mcintosh	31/3-023	74934,13+5584	Eric.Mcintosh@cern.ch
Central VMS Services	T.Whibley	513/R-033	74849	Tim.Whibley@cern.ch
AFS Service Manager	R.Többicke	31/1-024	78985	Afs.Support@cern.ch
Oracle (Relational Database)	S.Santiago	31/3-011	74134,13+5580	Oracle.Support@cern.ch
Distributed Printing	J.L. Vosdey	513/R-049	74976	Printer.Support@cern.ch
Magnetic Tape Operations	R.Minchin	513/R-009	13+5559	Richard.Minchin@cern.ch
Stager, tapes and TMS support				Tape.Support@cern.ch
CERN PROGRAM LIBRARY OFFICE				
Distribution of CERNLIB material	L.Besson	513/1-008	74951	Heplib.Support@cern.ch
CERNLIB Librarian	I.Mclaren	513/1-009	75010	Heplib.Support@cern.ch
LHC++ Librarian	J.Shiers	513/1-002	74928	Heplib.Support@cern.ch
DESKTOP COMPUTING				
NICE (Networked PCs)	A.Pace	31/1-019	72897	Nice.Support@cern.ch
Macintoshes	K.Pluntke	513/R-045	75015/160023	Klaus.Pluntke@cern.ch
UNIX Workstations and X Terminals				
- Front Desk	C.Ball	31/1-015	73349	Workstation.Frontdesk@cern.ch
- Team Leader	P.Defert	31/1-026	73923	Workstation.Support@cern.ch
Contact Accounts (according to OS):				
HP 700 s/w: Hp.Support@cern.ch - SUN software: Sun.Support@cern.ch - Solaris s/w: Solaris.Support@cern.ch				
RS/6000 s/w: Aix.Support@cern.ch - DEC's OSF/1 s/w: Osf.Support@cern.ch - X terminals: Xterminal.Support@cern.ch				
COMMUNICATIONS AND NETWORKS				
General network operational problems		513/R-034	74927	Network.Support@cern.ch
The complete and up-to-date list of contacts is at URL: http://network.cern.ch/wwwcs/public/contacts.html				
OTHER SERVICES				
Computer Security	J.Gamble	31/3-030	160261	Computer.Security@cern.ch
Tape Allocation to Experiments	H.Klein	13/3-024	72124,72060	Hansjorg.Klein@cern.ch
Web Office	M.Dimou	53/R-010	73356	WWW.Support@cern.ch
Software Development Tools	A.Khodabandeh	1/R-005	79099	swtool@ptsun00.cern.ch
Computer Science Library (mornings only)	J.Megies	513/1-024	72379	Jutta.Megies@cern.ch

SERVICE URLs (Some Useful Web Pages)

IT Home Page	http://wwwinfo.cern.ch/
Web Office	http://www.cern.ch/WebOffice
User Support	http://wwwinfo.cern.ch/support/
User Consultancy Office (UCO)	http://wwwinfo.cern.ch/uco
User Registration	http://wwwinfo.cern.ch/uco/user-registration
COmputing Books Selling (COBS) Service	http://wwwinfo.cern.ch/support/cobs-service
Computing Books Catalogue	http://consult.cern.ch/books
Accounting Service	http://wwwinfo.cern.ch/accounting/
CERN SP2 Service Page	http://wwwsp2.cern.ch/
NICE Home Page	http://nicewww.cern.ch/
UNIX Workstation and X Terminals Support	http://wsspinfo.cern.ch/file/
Software Development Tools	http://www.cern.ch/PTTOOLS/
Network Connection Request Form	http://network.cern.ch/register
TvScreen (Systems Status)	http://wwwcio.cern.ch/Others/TvScreen.html

Other Computing Newsletters

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

Available from M.Tate-Lavergne 513/2-010 72863 Monique.Tate-Lavergne@cern.ch

Editor L.Pregernig 31/1-006 72666 Ludwig.Pregernig@cern.ch

SEMINARS AND PRESENTATIONS

Designation	Usual Location	Organizer	Div.	Telephone	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	Gunter.Folger@cern.ch
		J.P.Porte	ECP	73457	Jean-Pierre.Porte@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	IT		IT-Training.Service@cern.ch
		R.Ramos	IT		

COMPUTER RESOURCES ALLOCATION AND DIVISIONAL CONTACTS

COMPUTER RESOURCES ALLOCATION COMMITTEE (COCOTIME):

Chairman: L.Foa/DG 4/1-048 78126 Lorenzo.Foa@cern.ch

Secretary: A.Ball/ECP 14/6-023 73849 Alan.Ball@cern.ch

DIVISIONAL REPRESENTATIVES FOR COMPUTING

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

Div.	Representative	Div.	Representative	Div.	Representative	Div.	Representative
AS	M. Draper	AT	P. Heymans	DSU	M. Draper	ECP	A. Grant
EST	Ch. Delamare	FI	M. Saitta	IT	A. Koppányi	PE	D. Duret
PPE	H. Klein	PS	J. Riche	SL	A. Hilaire	ST	E. Sanchez-Corral
TH	P. Nason	TIS	G.R. Stevenson				

NICE DIVISIONAL REPRESENTATIVES

An up-to-date list can be found at 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 at URL: <http://network.cern.ch/wwwcs/public/divreps.html>

Letters from the Editor

I take the opportunity of this “Letter from the Editor” to welcome the two new groups coming from the old ECP division, and which are part of IT since January 1st 1998. Both groups will continue with their former missions:

- The CO (COntrols) Group mission is to provide and maintain controls systems (also known as Slow Control, Industrial Control) to users in the Research Sector.
- The IPT (Information and Programming Technology) Group was formed to help the teams that produce software for and within CERN experiments to improve software development productivity and quality, by providing tools and methods. CNL readers are already familiar with this group which has provided, almost from its creation, a regular article for the CNL on “Software Development Tools Service”. This article always appears at the end of the “Desktop Computing” chapter.

As you will see, this first Computer NewsLetter of the Year 1998 contains lots of information, and all of the main chapters have been filled with some contributions.

One important item in this CNL is that the central computing facilities provided by the CERNSP service have been upgraded, at the beginning of March, to the new RSPLUS Public Login Service, based on PowerPC workstations.

The chapter on “Mail and News Issues” is now a regular section of the CNL, and will probably remain so for some time. The article on “Netscape 4 Mail”, which gives a series of practical tips and useful information, will probably interest many users, as there are already 700 users of this mail agent at CERN.

Regarding “Desktop Computing”, you will see *Changements des Droits d’Accès aux Répertoires “Home” de NICE / Changes of the Default Access Rights to NICE Home Directories*, that there will be (end of March) a change of policy concerning public access to files in home directories on the NICE servers. You should also read the article “ASIS News” to see the recent changes in the ASIS software repository, and send your feedback to the ASIS team.

In a long article on “Computer Networking at CERN” the “Communications Systems” group is asking users to help them improving the network performance, and explain what users should do (or should not do).

Finally, all of you who are interested in Web technology, and especially editing on the Web, should read the very detailed article on “HTML4 and Extensible Languages: the Web Prepares for the Future” in the “Text Processing” Chapter.

I hope you will enjoy reading all of this NewsLetter 230!

Nicole Crémel
CNL editor, IT/User Support

1. General

1.1 The 1998 CERN School of Computing

Carlo Vandoni, IT/DI

Date and place

The twenty-first CERN School of Computing will be held at the Hotel Savoy in Funchal, Madeira, Portugal, from Sunday 6 September to Saturday 19 September 1998.

Participation

The School is open to postgraduate students and research workers with a few years' experience in elementary Particle Physics, in computing or in related fields.

All other details of logistics, e.g. concerning accommodation, language, and travel can be found on the Web, at the URL <http://www.cern.ch/CSC/>.

Scientific programme

The School is based on the presentation of approximately 42 lectures and on related practical exercises on PCs or workstations, and the programme of the Schools is organised around four themes:

- Agent and Distributed Computing Technology.
- Intelligent Monitoring and Control.
- Petabyte Storage (data bases).
- Software evolution.

Agent and Distributed Computing Technology

This series of lectures will be complemented by exercises

This track explores the general field of distributed computing with a special emphasis on high-performance on the one hand, and on the use of agent technologies on the other hand.

The track is composed of two distinct parts. During the first week, background knowledge on high-performance distributed computing is provided, ranging from the statement of the problems to be resolved to the description of underlying technologies, including the use of toolkits such as DCE/CORBA. Case-studies will conclude this first part.

During the second week, the focus is on the use of agents written in Java. This method is applied to the specific field of Distributed Physics Analysis. After an introduction to the Agent technology, the class will discuss the application of the approach to Physics, including the use of submission and compute servers, and algorithm agents. Exercises where student write physics analysis algorithms in Java, and agent-based job submission systems, and finally merge their outcome into a global system, will complement the lectures.

Lecturers

- Mark Dönszelman, CERN, Geneva, Switzerland;
- Ian Foster, Argonne National Laboratory, USA.

Intelligent Monitoring and Control

This track discusses the application of artificial intelligence techniques to the monitoring and control of large experimental systems in High Energy Physics. The complexity and the expected long lifetime of the next generation of HEP experiments, in particular those that will take place at CERN's Large Hadron Collider, will place major challenges to the reliable operation of the detectors. Sophisticated data acquisition, monitoring and detector control systems will be developed to perform these tasks. Artificial intelligence techniques will play an important role in this context. Diagnosis, alarm handling and user assistance tools extracting information from a common knowledge base are examples of the facilities that can help to solve problems when the human experts are not available.

Lecturers

- Amilcar Cardoso, University of Coimbra, Portugal;
- Pere Mato, CERN, Geneva, Switzerland.

Petabyte Storage

Although High Energy Physics experiments have long generated data volumes far in excess of the usual norm, a number of experiments due to start in around 2000, and those planned to take data at CERN's Large Hadron Collider, will push these limits even further. The LHC experiments are expected to take some 5 PetaBytes (PB) ($1\text{ PB} = 10^{15}$ bytes) of data per year, giving rise to a total data sample, integrated over 15-20 years' running, of perhaps as much as 100 PB. In addition to the sheer volumes of data involved, the extremely long time-scales — by the standards of the IT industry or the career of an individual physicist — mean that great care must be taken to protect against change. Thus, the use of standard, and, if possible, commodity solutions are necessary.

This track discusses the basic components that are being investigated in order to provide data storage and data management solutions to future HEP experiments. Not only will the standards that are involved be discussed, but so too will concrete solutions that have already been used in production for storing and managing physics data. The main components that are currently being studied include the High Performance Storage System (HPSS), a mass storage system built according to the IEEE Computer Society's reference model for such systems, and Objectivity/DB, an ODMG-compliant

Object Database Management System (ODBMS). The track will cover these two components, plus the relevant standards and architectural choices, as well as the CERN RD45 project which is investigating these components as a solution to handling LHC event data.

Lecturers

- Dick Watson, Lawrence Livermore National Laboratory, Livermore, USA;
- Drew Wade, Objectivity Inc., Mountain View, USA;
- Dirk Düllmann, CERN, Geneva, Switzerland.

Software evolution

Smooth software evolution requires a well established development process. Starting from the standard software process model (CMM) of the Software Engineering Institute, the lecturers will explain the activities required to start improving the process and expand further on configuration management and quality through software metrics. For the exercises, students would be requested to work on their own projects, applying software metrics to the source code and organizing all the development documents with a version control system (CVS - *Concurrent Version System*).

Lecturers

- Klaus Ackerstaff, CERN, Geneva, Switzerland;
- Alberto Aimar, CERN, Geneva, Switzerland;
- Gottfried Kellner, CERN, Geneva, Switzerland;
- Arash Khodabandeh, CERN, Geneva, Switzerland;
- Bob Jacobsen, University of California, Berkeley, USA.

How to apply

If you have a WWW viewer that supports forms then you may fill out and submit the CSC '98 WWW application form at the URL <http://www.cern.ch/CSC/CSC98-4.html>. If

you decide to use the Web to send in your application form do not forget that you also need to send *A Formal Letter of Reference*.

If you choose not to use the Web, then you must fill out a printed copy of the *plain text version of the application form* available at the URL <http://www.cern.ch/CSC/AppForm.txt>.

When applying to attend the School, each student is requested to provide a summary in English of about 100 words in length, describing her/his current work. Applicants are requested to forward their summary to Miss J. Turner by e-mail (Computing.School@cern.ch).

Candidates should forward the completed application form and formal letter of reference to Miss J. Turner as specified in the Section "Enquiries and Correspondence".

The deadline for receipt of applications is 15 May 1998.

For more details concerning the application cost and information concerning financial support or selection are given on the Web at URL:

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

Enquiries and Correspondence

All enquiries and correspondence related to the School should be addressed to:

Miss J. Turner
CERN School of Computing
CERN
CH-1211 GENEVA 23
Switzerland

E-mail	Computing.School@cern.ch
Telephone	+41 22 767.50.49
Telefax	+41 22 767.71.55

1.2 The RSPLUS Public Login Service

Tony Cass and Pietro Martucci, IT/PDP

The new RSPLUS interactive service was introduced on March 2nd, replacing CERNSP. As March 2nd was also the deadline for this CNL, we cannot be sure about the progress for this article but we are hopeful that the service will have been well received, since very few problems were reported during the trial period.

The only important difference that early RSPLUS users have noticed is that the `lpr` command cannot be used to print files. Although `lpr` was made to work on the SP2 when the CERNSP service was first introduced, this command has never been fully supported on the other PLUS services, whilst `xprint` has been the supported printing command on

UNIX services for some time. There have been many improvements to `xprint` since the CERNSP service was first introduced and, in particular, it is now possible to pipe output to `xprint`. As such, we are sure that users will find that `xprint` is an adequate replacement for `lpr`, although we understand the inconvenience this change has caused.

Overall, however, we hope that users appreciate the migration to RSPLUS as the new service. Using 15 twin-processor PowerPC workstations it should provide much better performance than the 24 single processor SP2 nodes.

For batch users, RSPLUS supports job submission with the

llsubmit command and general *LoadLeveler* query commands also work. However, the special CERN extension commands (such as llmore and llps) cannot be used from RSPLUS as they rely on special SP2-based authentication. If you need to use these commands then it is still possible to connect to the SP2 nodes by logging into the sp2pub service rather than rsplus (or cernsp). However, as detailed in another article in this CNL, *LoadLeveler* will be replaced by LSF as the job scheduler for the public batch service and the RSPLUS nodes have a full interface to the LSF batch environment.

Phasing out the SP2-based CERNSP service

Although the SP2 nodes that supported the CERNSP service can still be accessed by using the sp2pub alias, we will be

removing sixteen of the nodes from the cluster at the end of March. Further nodes may be removed from the sp2pub cluster at a later date depending on the load. The sp2pub service will remain available until the end of the *LoadLeveler*-based batch service but will be turned off completely once the migration to LSF is complete, expected to be at the end of June.

As for the cernsp name, this can still be used but we strongly recommend people to use the rsplus name for the new interactive service. The cernsp name will be removed from the list of central services on NCD X terminals and for PC network connections on June 1st. The CERNSP alias will be removed completely on September 1st.

1.3 The Public Batch Service: An Upgrade and the Migration to LSF

Tony Cass, Pietro Martucci and Alessandro Miotto, IT/PDP

With the end of maintenance for most of the SP2 nodes, the public batch service will be entirely based on IBM PowerPC workstations. 18 PowerPC nodes are in use today and a further 15 systems were purchased at the end of last year. These new machines are equipped with 332 MHz PowerPC 604e CPUs, each rated at 100 CERN Units and thus provide a significant increase in capacity. The total PowerPC based capacity will be 2040 CERN Units compared to 990 CERN Units for the current configuration (and 450 units for the original SP2 batch service in 1995).

However, as we move away from the SP2, we will also be changing to use LSF as the batch scheduler rather than *LoadLeveler*. LSF is already in wide use for SHIFT batch services and will replace both NQS and *LoadLeveler* to become the single batch scheduling system at CERN by the end of 1998.

LSF has already been installed on the 15 new machines and on the RSPLUS nodes. Adventurous users can look at the existing LSF Web pages (URL <http://wwwinfo.cern.ch/pdp/lsf>) and try submitting LSF jobs now. We will, however, be providing more documentation in the coming weeks. A brief introduction to the LSF queues and a sample LSF job are included below.

The key step for the introduction of the LSF based public batch service will be the introduction of CPU quotas based on the COCOTIME allocations, expected to be at the beginning of April. Following this we will start to migrate the existing PowerPC nodes to LSF and leave only the SP2 nodes running *LoadLeveler*. The end of the *LoadLeveler* batch service is currently scheduled for June 28th.

LSF queues

The batch queues on virtually all of the batch clusters running LSF are defined with an indication of the queue length in "Normalised" (or "New") CERN Units. This "NCU" is simply scaled by a factor of 100 from the existing CERN Unit definition. Conveniently, therefore, the queue definitions for the new Public Batch machines show how much real CPU time a job can utilise. The 4 queues defined on the new cluster are as follows:

8nm	8 New Minutes, equivalent to about 13.5 CERN Unit hours
1nh	1 New Hour, equivalent to 100 CERN Unit hours
8nh	8 New Hours, equivalent to 800 CERN Unit hours
1nd	1 New Day, equivalent to 2400 CERN Unit hours

We will be reviewing this set of queues in the light of experience.

A Sample LSF Job

Finally, we include here a sample script ("example2" from the LSF examples at the URL <http://wwwinfo.cern.ch/pdp/lsf>) that will run a Fortran job under LSF. If this script is saved as myjob, the command to submit and run the job is simply `bsub < myjob`.

As you can see, LSF directives are given in lines starting #BSUB, e.g. #BSUB -J example2 or #BSUB -q 8nm, where the options q and J are just those that can be given on the bsub command line. A full list of bsub options is available in the bsub man page (type `man bsub`).

```
#!/usr/local/bin/zsh
#
# As usual, the first line defines the shell
```

```
# Lines beginning '#BSUB' are LSF directives
#
# Give the job a name
#
#BSUB -J example2
#
# We don't need much time so the 8 normalised
# minute queue is fine
#
#BSUB -q 8nm
#
# Copy the in-line Fortran to a work file.
#
cat > temp.f <<EOF
PROGRAM TIME
Do M = 1,100000
  If ( 2*(k/2) .eq. k ) k = 0
  Do I = 1,100
    Do J = 1,100
```

```
      K = K + I*J
    Enddo
  Enddo
Enddo
call timel(tleft)
call timex(tused)
write (6,'(The answer:,I9)') k
write (6,'(Tused =,F9.2/Tleft =,F9.2)') tused,tleft
END
EOF
#
# Compile and link with CERNLIB
#
hepf77 temp.f 'cernlib'
#
# Execute the program
#
a.out
```

1.4 Migration of 3480s to Redwood: Progress Satisfactory

Charles Curran, IT/PDP

At the time that the introduction of the new STK automation equipment was announced, users were informed that IT intended to try to move user data off 3480 media onto the newer and fully automated *Redwood* media.

The 3480 cartridge was still in frequent use, but no longer supported by automation at CERN. The basic proposal was to copy all 3480 cartridges in the Tape Vault (some 200,000) onto 25 GBytes *Redwood* volumes, using VID-mapping features of the Tape Management System (TMS) to make this as transparent as possible to users. The new copies would be accessible for READ only, but this was not expected to be a great difficulty as this mode corresponded to the great majority of usage. The new media was to be provided by IT, as was the actual copying scheme and manpower.

I am pleased to say that this project, thanks to the goodwill and cooperation of all concerned, is now in its closing stages. It took a little longer than expected to start up, but

from late July 1997 it has been possible to keep eight copy streams active for almost 168 hours/week, using the Tape Vault self-service copy stations and some 170 GBytes of dedicated stage disk. When no problems were met, over 5000 volumes could be handled in a week. You can find more details at the URL:

<http://wwwinfo.cern.ch/pdp/vm/migration.html>

As the 3480 works finishes, work has started to migrate 3490s in the same way, starting with the busiest volumes and those that were previously in the SMCG robot. Although these volumes are not so numerous (some 20,000) they take longer to copy: they contain more data, and we can expect to use only 6 streams at best. Nevertheless, it is still expected that all tape handling apart from preparation and import/export support activities will be automated by the end of 1998.

1.5 The Future of VXCERN

Judy Richards, IT/DIS

The future of VXCERN has been discussed in the last two FOCUS meetings in October and December 1997. It is clear that certain parts of the existing service must be kept running until the end of LEP data taking, in order to support activities which are dependent on VMS. However, a rather large percentage of current VMS users are using the service almost exclusively for electronic mail and news reading. These users should be encouraged to move to other alternative services.

Parts of the service for which maintenance is difficult should be stopped.

Thus, following the recommendations of FOCUS, we will

- stop registration of new VXCERN accounts on April 1st,
- stop the tape service and SHIFT disk access on July 31st,

- stop archiving of files on April 1st. Note that archive files will be available for restore until 2001. Users are responsible for transferring any files they want to keep to other systems.
- Since Digital will provide “Year 2000” fixes for VMS 6.2 we will not upgrade the operating system to VMS 7.

In order to facilitate the migration of e-mail to the central MailServer, the VMS version of *Pine* has been installed. See the article “*Pine Now Available on VXCERN*” on page 9 for more information.

A review of the news reading facilities on UNIX and NICE has started with the objective of improving them to provide the more popular features of VAXNEWS.

In December 1998 a major review of VXCERN accounts will be made with the experiments in order to identify those accounts which should be kept open beyond January 1st 1999.

Please note that these changes to the VXCERN service do not imply a change of policy for the maintenance of the VMS version of CERNLIB: new versions will still be available via the ASIS service.

1.6 WebOffice Tools

Maria Dimou-Zacharova, IT/WO

There are some tools available on our central Web server to help users check or improve the quality of their HTML pages.

- <http://www.w3.org/Security/Faq/www-security-faq.html>
Web Security FAQ (link available from the top of the *Web Office home page* at the URL <http://www.cern.ch/WebOffice/>).
- <http://www.cern.ch/WebOffice/Search>
Explains how to make sub-searches of your data indexed by the CERN Intranet search engine.
- <http://www.cern.ch/WebOffice/Projects/Styles>
Provides a set of tools for: HTML syntax check, spelling check, link validation check and introduction of META-data.
- <http://www.cern.ch/WebOffice/Statistics/pagesize>

Counts HTML and image size for any URL(s) you give.

- <http://www.cern.ch/WebOffice/DoctorHTML>
Analysis tool for link validation check.
- <http://www.cern.ch/WebOffice/Courses/>
Contains useful information on HTML, Cascading Style Sheets, XML, Frames, clickable images, JavaScript with examples.
- <http://www.cern.ch/WebOffice/Projects/Calendar>
Contains requirements and products for a shared calendar service.
- <http://www.cern.ch/WebOffice/guides.html>
Contains some pages of the W3 Consortium with useful FAQ and Web editing/conversion tools.

If you have some specific questions, please feel free to contact us at the Web Office (Web.Support@cern.ch).

1.7 New User Registration Form

Miguel Marquina, IT/User Support

As a byproduct of the rationalization of the CERN's Human Resource Database (HR) and the IT Computer Centre Data Base (CCDB), we have revised the User Registration form to be filled in for new users of Central Computing facilities.

This form is available on paper at the UCO premises and, more conveniently, electronically at the URL <http://consult.cern.ch/uco/user-registration>.

The main changes are:

- There are no longer two separate forms (“green” and “yellow”) to distinguish between “General Users” and “Outside Users”. The new form is white, thus compatible with your own printed copies or the faxed forms.
- There is additional mandatory information which has to be filled in, in order to ensure correct matching with the registration in HR. Please note, as indicated in the form, that such information will be treated confidentially and only for the purpose stated above.

2. Mail and News Issues

2.1 Netscape 4 Mail

Arnaud Taddei, IT/DIS

Netscape 4 Mail was introduced at CERN in June 1997 on the Macintosh platform and later on NICE and on UNIX systems. Users started to require it for various reasons:

- need to access their mail from different platforms, particularly from home;
- need to use other features of *Netscape 4* like Java;
- need to move from other obsolete mail systems;
- already happily running version 3 of *Netscape*.

We therefore started to organise its support (documentation, debugging, etc.). Look at URL <http://wwwinfo.cern.ch/service/mail/clients/netscape>. There are at present around 700 users of this product at CERN.

The general consensus is that *Netscape 4 Mail* is usable but not of the quality we had hoped. It offers more features than version 3, in particular supporting access to your mailbox from multiple locations using IMAP.

The product has quite a number of more or less serious bugs. The company has agreed to a certain number of changes promising a better product with *Netscape 5* some time in 1998. A complete list of problems can be found at URL <http://wwwinfo.cern.ch/service/mail/problems/netscape>.

In order to alleviate the effects of one of the most serious problems, we have managed to find a workaround by tuning one of our Mail Servers and implementing a new mailbox format which offers better connectivity for users, in particular the ones using *Netscape 4 Mail*.

In addition we have produced some explanations and recommendations that should help you get a better service with *Netscape*. They are summarized below and the full details can be found at URL

<http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#details>

- **Differences between IMAP and POP** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#reminder>) The Mail Server offers location independent access to your mailbox and folders using the IMAP protocol.

- **Creation of folders on the Mail Server** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#creation>) How to avoid the main traps with *Netscape* when creating folders on the Mail Server?
- **Moving your AFS folders to the Mail Server** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#moving>) How to move all your folders to the Mail Server: in particular, your AFS folders? If you are not just using UNIX/AFS systems this will help you a lot!
- **Delete and “expunge” mechanism** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#delete>) How to avoid serious problems by using the recommended “delete” and “expunge” mechanisms with *Netscape*?
- **Checking for New Mail** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#rate>) How to avoid problems caused by having a polling rate which is too high in *Netscape 4 Mail*?
- **Have your account on the right Mail Server** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#mail3>) How to improve your life with *Netscape* by getting your Mail Server account moved to the right server?
- **A new mailbox format** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#mbx>) What is this new format on the Mail Server which really helps a lot when using *Netscape 4 Mail* and what do I need to understand about it?
- **Netscape 4 Mail on UNIX** (URL <http://wwwinfo.cern.ch/service/mail/clients/netscape-recommendations#unix>) Are the UNIX versions of *Netscape 4 Mail* usable? (It depends on the platform.)

2.2 Unsolicited Electronic Mail

Judy Richards, IT/DIS

Since November last year we have been actively trying to reduce the amount of unsolicited electronic mail, commonly known as SPAM mail, that arrives in CERN mailboxes. Unfortunately this is not an easy job since the spammers are constantly finding new ways of breaking our defences! Gone

are the days when you could block spam mail with a simple filter on the ‘From:’ address of a sender. Today spammers are “forging” addresses in the headers and envelopes of mail in such a way that it is difficult to distinguish between legitimate mail and spams.

An additional problem, that does not directly affect most CERN users but does use CERN resources and affects the reputation of CERN, is caused by the spammers who try to make mail look as if it comes from a CERN machine. They either try to route the mail via a CERN machine or simply forge the mail headers to a CERN address so it will have less chance of being blocked by the recipient's anti-spam defences.

Although some anti-spam measures can be automated, for most someone has to look at the system logs and spam reports from users and make a human judgement as to what should be blocked. This of course takes time and so there is a limit as to what we can do. In addition it is difficult to measure how effective our anti-spam measures are. Although we know how much mail we are rejecting at the CERN gateways, each of these could be "fanned out" to multiple recipients at CERN, but it is prohibitively resource consuming to find out how many. And of course we have no way to measure how much gets through the filters into your personal mail boxes.

Since it is difficult to distinguish between spam mail and genuine mail, "innocent" people may occasionally be affected by our anti-spam measures. These are summarized below so that you are aware of what we are doing.

Banned domains

Mail from all userids in these domains are banned. A domain is added to this list when we have seen spams from multiple users of this domain. This ban of course also affects "innocent" users of the domain.

Banned IP addresses

Filtering is done on the physical IP address. This catches spammers who are forging or keep changing the name of their machine.

Banned senders

This list is useful only against the most amateur of

spammers since most now use a different 'From:' address for every spam.

A list of domains for which mail is accepted only if the sender and relay nodes are in the same domain

In this list we find some of the larger Internet providers whose addresses are often forged by spammers. This filter can however affect some "innocent" users. For example, if a user with the address 'user1@aol.com' sends mail to 'user2@in2p3.fr' and 'user2' has an automatic forward to 'user2@cern.ch', the mail will be rejected since the 'From:' address is 'aol.com', but the relay address is 'in2p3.fr'. Automatic forwarding from within CERN is not affected.

Attempts to use CERN as a relay are rejected

Any mail for which neither the sender address nor the recipient address is registered in the CERN name servers is rejected. This can affect two groups of CERN users:

1. If you are at home or travelling you will not be able to send mail to someone outside CERN using a CERN machine as your SMTP server. The standard CERN configurations for mail clients set the SMTP server to 'smtp.cern.ch'. If you are outside CERN you must change this to the SMTP server of your Internet provider or the lab where you are physically.
2. If you try to send mail from a machine that is not correctly registered in the CERN central name servers it will be rejected.

A list of the domains that are currently banned can be found at URL <http://wwwinfo.cern.ch/service/mail/problems/SpamDomains>. This list is updated once per day and today contains only the "banned domains". When the next version of the software is introduced it will also contain addresses falling into the other categories.

2.3 Macintosh Mail Migration

Mats Møller, AS/SU

At the *DeskTop Forum* in the meeting of February 19th, it was decided to restart the migration away from *QuickMail*. After careful testing and user evaluation the following strategy will be adopted:

- The majority of the users will be migrated to *Eudora Pro (POP3)*, for which we will acquire a site license. *Eudora* has been greatly appreciated by the group of pilot users. It is considered to be very easy to use and

has minimal memory requirements (i.e. it runs well on all Macintosh models used at CERN).

- Users needing IMAP for mobility reasons will be proposed *Netscape 4*.

The migration will be carried out on a Mail Center by Mail Center basis and each user will receive further information in due course.

2.4 Pine Now Available on VXCERN

Tami Kramer, EP/LE

Pine is a simple screen-based mail and news program which runs on several platforms. At CERN it is available on Unix, NICE and now VXCERN. Working in conjunction with an IMAP server, it allows users to attach to two or more mailboxes in the same session, and move mail folders between them. *Pine* has been installed on VXCERN to aid VMS mail users in migration to the MailServer.

Some documentation on using *Pine* can be found at the URL <http://www.washington.edu/pine/> although the program is pretty much self-contained.

Configuring *Pine*

In order to use *Pine* you have to customize it with your personal information, such that you can view both your VMS Mail mailbox and your MailServer mailbox from the same interface. This will allow you to move folders to the MailServer, preserving the integrity of your mail messages. This procedure is meant as a migration tool, rather than a standard way of reading your mail. You will be prompted for your MailServer username and password when you attempt to open your MailServer mailbox.

The initial recommended setup is achieved by executing the command `PINESETUP`, which will:

- install a file called `.PINERC` in your VMS home directory;
- place your MailServer loginid (by default your VMS loginid) in it;
- point your incoming mailbox to be that on the MailServer.

Your VMS mail is assumed to be in the directory `SYS$LOGIN:[MAIL]`, but you can change it through *pine* via the command sequence:

M (for Main Menu)

S (for Setup)
C (for Configuration)

then find the “VMS Folders” definition in the “folder-collections” option:

```
folder-collections =
  Mail Server Folders {loginid.mailbox.cern.ch}mail/[]
  VMS Folders ~/mail/[]
```

The file syntax is UNIX (as this is a VMS port of a UNIX application). If you change it to:

```
VMS Folders ~/[]      or,
VMS Folders ~/anotherdir/[]
```

you will access instead all mail folders in the `MAIL.MAI` file in your VMS home directory or `SYS$LOGIN:[ANOTHERDIR]` respectively.

If all goes fine, your initial setup would look like this:

```
personal-name      =
  <No Value Set: using ... >
user-domain        =
  <No Value Set>
smtp-server        =
  <Value is Fixed: using smtp.cern.ch>
nntp-server        = news.cern.ch
inbox-path         = inbox
news-collections   =
  <No Value Set: using *{news.cern.ch/nntp}[]>
folder-collections =
  Mail Server Folders {loginid.mailbox.cern.ch}mail/[]
  VMS Folders ~/mail/[]
```

You may modify it at any time using the relevant *Pine* commands for configuration.

3. Desktop Computing

3.1 Changements des Droits d'Accès aux Répertoires *Home* de NICE¹

Alberto Pace, IT/DIS (pour l'équipe NICE)

Aujourd'hui la majorité des répertoires *home* sur les serveurs NICE (SRVx_HOME, x=1,2...7) sont par défaut lisibles par tous les utilisateurs. Comme cela a créé un certain nombre de malentendus dans le passé, les droits d'accès sur les répertoires 'home' seront changés le samedi 28 Mars 1998 pour interdire l'accès en lecture par d'autres utilisateurs. Ceci mettra la politique du service NICE en accord avec les autres services centraux de la division IT.

Un nouveau sous-répertoire appelé *public* sera créé dans chaque répertoire 'home'. Les fichiers sauvegardés dans ce répertoire seront lisibles par défaut par tout le monde (mais pourront être modifiés seulement par le propriétaire, comme d'habitude). Tous les sous-répertoires créés au-dessous du répertoire *public* seront également lisibles par d'autres utilisateurs. Tous les fichiers ou sous-répertoires créés dans le répertoire 'home' qui ne sont pas sous le répertoire *public*, ou le sous-répertoire de *www* (voir suite), seront invisibles à d'autres utilisateurs.

Les droits d'accès explicites donnés par des utilisateurs à d'autres utilisateurs ou groupes demeureront sans changement.

Première action nécessaire – si vous partagez des fichiers avec d'autres utilisateurs

Tous les fichiers dans votre répertoire "home" qui doivent être consultés par d'autres utilisateurs doivent être déplacés dans votre répertoire *public* avant le samedi 28 Mars 1998. A partir de cette date, les fichiers à l'extérieur de *public* ou *www* (voir suite) ne seront plus accessibles par d'autres utilisateurs jusqu'à ce qu'ils soient déplacés.

Les fichiers qui sont dans le répertoire *private* resteront cachés.

Deuxième action nécessaire - si vous utilisez *niceworld.cern.ch*

Aujourd'hui, les fichiers lisibles publiquement situés dans votre répertoire 'home' sont accessibles sur le Web avec

le service de *niceworld.cern.ch*. L'adresse <http://niceworld.cern.ch/~loginame/path/filename> affiche le fichier `\\srvx_home\usr_xx\home\loginame\path\filename` du répertoire 'home' de l'utilisateur.

Après le samedi 28 Mars 1998, un nouveau sous-répertoire appelé *www* sera créé dans chaque répertoire 'home' utilisateur. Les fichiers sauvegardés dans ce répertoire seront lisibles publiquement et la syntaxe <http://niceworld.cern.ch/~loginame/path/filename> affichera le fichier `\\path\filename` du sous-répertoire *www* (`\\srvx_home\usr_xx\home\loginame\www`). Les utilisateurs pourront créer une page d'accueil (*Home Page*) avec comme nom `\\srvx_home\usr_xx\home\loginame\www\default.htm` qui sera affichée en accédant à l'URL <http://niceworld.cern.ch/~loginame>.

Comme auparavant, les utilisateurs peuvent également permettre la navigation à l'intérieur de leur sous-répertoire `\\srvx_home\usr_xx\home\loginame\www` en créant un fichier *niceworld.ini* contenant les lignes:

```
[NiceWWW]
DirectoryBrowsingEnabled=1
```

Dans tous les cas, l'accès au Web sera limité au sous-répertoire *www* du répertoire 'home' de l'utilisateur.

Note: Pour votre convenance, si un fichier nommé *default.htm* existe dans la racine de votre répertoire "home" et si un fichier avec le même nom n'existe pas dans le sous-répertoire *www*, le fichier *default.htm* sera automatiquement copié dans le sous-répertoire *www*. Une édition manuelle de ce fichier sera nécessaire pour corriger certains liens sur des images incluses dans ce document.

La Figure 3.1 montre la situation par défaut avant et après le samedi 28 Mars 1998. Notez que ce changement a été discuté au *Desktop Forum* lors de la réunion du 4 décembre 1997. Voir le document à l'URL <http://www.cern.ch/Official/Desktop/DF17.html> ou contacter votre représentant du *Desktop Forum* pour de plus amples informations.

3.2 Changes of the Default Access Rights to NICE Home Directories²

Today the majority of home directories on the NICE servers (SRVx_HOME, x=1,2,...7) are readable by everyone by default. As this has created some misunderstanding in the past, on Saturday March 28th 1998, the file access protection on users' home directories will be changed to forbid open

access by other users. This will bring the NICE service into line with other central service policies.

A new subdirectory called *public* will be created in every user home directory. Files saved in this directory will be pub-

¹ See Section 3.2 for the English version of this article.

² Voir la section 3.1 pour une version française de cet article.




Figure 3.1: Situation par défaut *avant* (en haut) et *après* (en bas) le samedi 28 Mars 1998

lically readable by default (but writable only by the owner, as usual). All subdirectories created below the public directory will also be publically readable. All subdirectories created in the home directory that are not under the public or the www subdirectory (see later) will be hidden to other users.

Explicit access rights assigned by individual users to other individual users or groups will remain unchanged.

First necessary action - if you share files with other users

All files in your home directory that should be accessed by other users *must* be moved into your public folder before Saturday March 28th 1998. From that date, files outside the public or the www (see later) folders won't be accessible for read to other users until they are moved into the public directory.

Files that are in the Private directory will remain hidden.

Second necessary action - if you use nicewww.cern.ch

Today, all public readable files located in your home directory are accessible from the Web using the

nicewww.cern.ch service. The syntax `http://nicewww.cern.ch/~loginame/path/filename` returns the file `\\srvx_home\usr_xx\home\loginame\path\filename` from the user's home directory.

After Saturday March 28th 1998, a new subdirectory called `www` will be created in every user home directory. Files saved in this directory will be public readable and the syntax `http://nicewww.cern.ch/~loginame/path/filename` will return the file from the user's home directory `\\srvx_home\usr_xx\home\loginame\www\path\filename`.

Users can define a home page by creating a file named `\\srvx_home\usr_xx\home\loginame\www\default.htm`, it will be returned with a reference of the form:
`http://nicewww.cern.ch/~loginame`.

As before, users can also enable directory browsing by creating a file `\\srvx_home\usr_xx\home\loginame\www\nicewww.ini` containing the lines:

```
[NiceWWW]
DirectoryBrowsingEnabled=1
```

But, in all cases, browsing from the Web *will be limited to* the `www` subdirectory of the user's home directory.

Note: For your convenience, if a file `default.htm` exists




Figure 3.2: Default situation *before* (top) and *after* (bottom) Saturday March 28th 1998

in the root of your home directory and if a file with the same name does not exist in the `www` subdirectory, the file `default.htm` will be moved down into the `www` subdirectory. A manual editing of this file may be necessary to adjust links to embedded images.

Figure 3.2 shows the default situation before and after Saturday March 28th 1998. Please note that this change was discussed at the *Desktop Forum* in the meeting of December 4th 1997. See the document at the URL <http://www.cern.ch/Official/Desktop/DF17.html> or contact your *Desktop Forum* divisional representative for further information.

3.3 Operating System Certification - Present List

Philippe Defert, IT/DIS

Over the coming months, here are our plans for certifying the various versions of the UNIX operating systems which we support on the desktops. For further information on any of these, or comments on the plans, please contact us via the URL <http://wwwinfo.cern.ch/dis/ose/Staff.html>.

Note that in a number of cases we are dependent on software coming from vendors. As the situation changes, appropriate news items will announce the latest status.

- **SUN Solaris Certified 2.5.1**
We will soon start working on 2.6, since AFS for 2.6 has just arrived.
- **HP-UX Certified 10.20**
HP released 10.20 (ACE). Transarc are currently certi-

fying AFS for this (but were reported that 10.20 AFS seems to run). They have no plans for a 10.30 port, since they claim that HP never put this on General Release.

HP 11 has just arrived. Transarc are starting the port, so expect something in 6-12 weeks.

We believe there is little interest in either 10.30 or 11 at present. (Are we wrong?)

- **D-UNIX Certified 4.0.b**
4.0.c is not certified but works (inc. AFS for release b) Digital has just released 4.0.d. Transarc is working on the port (expect more news in 2 weeks). We know that some users want this release, but until we hear from Transarc we are unable to give a date.

- **AIX Certified 4.1**
Certification of 4.2 about to be started.
We anticipate that some users may want 4.2 soon (and

may be 4.3 after that), but the 4.3 kit has only just arrived at CERN, and there is no news from Transarc on their plans for it.

3.4 Operating Systems Availability on ASIS and ASIS News

Philippe Defert, IT/DIS

The status of the ASIS application software repository is as follows:

- **sun4x_55:** SUN Solaris 2.5.1
- **hp_ux102:** HP/UX 10.20
- **rs_aix41:** IBM AIX 4.1.5
- **rs_aix42:** IBM AIX 4.2, (partial and unsupported)
- **alpha_dux40:** DEC OSF 4.0b
- **sgi_53:** SGI Irix 5.3 (not supported as desktop)
- **sgi_64:** SGI Irix 6.4 (partial, not supported as desktop)
- **i386_linux:** Linux redhat 4.1

At the end of March a few changes will be introduced in the configuration of some ASIS applications:

- **gv** replaces **ghostview** as the default front end for **ghostscript**, the PostScript viewer. The support for **ghostview** disappeared. **gv** is the natural follow-up of the **ghostview** project. **gv** is already available in ASIS. It has a 3D look and feel and understands pdf. The configuration of the applications which have to display PostScript files, such as WWW browsers, mail or news readers, will be changed to point to **gv** instead of **ghostview**.
- The support for **tk 4.2** and **tcl 7.6** has been dropped by the **tcl/tk** developers. The commands **tclsh** and **wish** will be updated to point to **tcl/tk 8** (8.0 or 8.1, whichever is the more stable version, to be installed at the end of March). **tclsh** version 7.6 will

remain available as **tclsh7.6** and **wish** version 4.2 as **wish4.2**. In general, version 8 is backward compatible with 7.6/4.2. Note that **tcl** version 7.3 and **wish** version 3.6 are still in ASIS but without any support **at all**: it is advised to migrate to version 8.0 as soon as possible.

- The **gcc** project now has two different development lines:

1. **gcc**, done by the FSF, which will handle very conservative but stable versions,
2. **egcs**, sponsored by Cygnus, which will include more innovative feature for C++ and optimisation.

gcc version 2.8.0 and **egcs** 1.0.1 are more or less at the same level. Unfortunately, **gcc** and **egcs** cannot be both installed with the same prefix. **egcs** contains **g77** but **gcc** 2.8.0 does not have any **g77** co-operating with it. Thus, on ASIS, the **gcc** version is still 2.7.2.3 and **egcs** was installed in the directory `/usr/local/egcs/bin`. As soon as a version of **g77** is available for **gcc** 2.8.0, the latter will replace the old version of **gcc** in the production area of ASIS. It is difficult to decide if **gcc** should be dropped for **egcs** or not!

In the future, it is planned to call for a *UNIX Users Forum* to decide on these matters. If you are interested in participating in such a forum, please contact the author (Philippe.Defert@cern.ch).

3.5 New ASIS Mailing Lists

Lionel Cons, IT/DIS

We have created a set of mailing lists linked to important ASIS products. The first lists are:

- **gcc**: for the GNU "compilers": **gcc**, **g77**, **gdb**...
- **perl**: for Perl and its main modules: **Tk**, **LWP**...
- **tk**: for **Tk/Tcl** and **tcl**
- **emacs**: for GNU **Emacs** and its extensions: **VM**, **AUCTEX**...

The full name of each list is `asis-<product>@listbox.cern.ch`, for instance `asis-perl@listbox.cern.ch`. To subscribe, use the Listbox service as documented at URL <http://wwwinfo.cern.ch/service/mail/listbox>.

These lists are intended to discuss freely and exchange information about these important ASIS products. All changes in

ASIS concerning these products (new version available for tests, migration planed, new relevant products, etc.) will be sent to the appropriate list(s). We expect that advanced users of these products will subscribe and that they will help to share expertise inside CERN. Note that these are not "support lists": there is no guarantee that someone on the list will be able to answer questions or solve problems. However, the relevant ASIS product maintainers will be on the lists.

You can browse through the HTML archives of these lists at URL <http://wwwinfo.cern.ch/~majordom/news/asis-<product>/>, for instance, <http://wwwinfo.cern.ch/~majordom/news/asis-gcc/>.

If needed, we can create other lists for other ASIS products. Send your suggestions to the email address asis.support@cern.ch.

3.6 Zephyr: The Single Notification System for All Users of Central Computing Services

Tony Cass, IT/PDP

The introduction of *Zephyr* for PC users at the end of January has led to a number of questions about the alarm messages that started to appear on PC screens. In part, these questions were prompted not just by the fact that these new messages appeared, but also because so many appeared in the first weeks after the introduction of the new service!

A notification system is needed to allow the operators of the Central Computer services to tell users of any problems that may affect them or their colleagues. Using *Zephyr* as the notification system allows us to send a message to almost all computer users. By sending such messages we can avoid having many users call the User Consultancy Office or the operators—or waste time trying to find out if they are doing something wrong. Fortunately, such incidents are relatively rare and no *Zephyr* warnings were sent out at all in the last 2 weeks of February, for example.

More services are now being provided in common for UNIX and PC users and *Zephyr* helps us to keep all users informed

of what is happening with central services. As such, we hope you can accept the occasional *Zephyr* message. These messages are only sent when necessary and one day such a message may prevent you from wasting time trying to solve a problem that lies elsewhere.

Finally, one or two users have complained that it is not just the Operators that can send messages, and that their work can be disturbed by other users. However, you can control the messages that *Zephyr* displays with a `zephyr.sub` file in your NOVELL home directory, i.e., a file “J:\zephyr.sub”. If this file contains the line

```
!message, personal, %me%
```

then other users will not be able to send you messages. You can create such a file yourself by starting *Notepad* or *Wordpad*, entering this line and saving the file in your home directory (J: drive).

3.7 Software Development Tools Service

Arash Khodabandeh, IT/IPT

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/PTT00LS/>

Below is the list of currently available tools (for the most recent news look at the URL <http://www.cern.ch/PTT00LS/SoftNew.html>).

Requirement Engineering

- **PSS-05 URD Template** (group product): For FrameMaker and WebMaker.

Analysis and Design

- **ADAMO** (group product): ER design and programming system.
- **OMW**: Martin/Odell analysis and design tool with C code generation.
- **Rose/C++**: OMT/Booch analysis and design tool with C++ code generation.
- **StP/OMT/Booch**: OMT/Booch 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 and Fortran. There now are four tokens for the SNiFF+/Fortran parser supporting F77 and F90.

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

- **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.
- **WebMaker**: Generic configurable converter from FrameMaker to HTML (for UNIX, Mac and Windows).
- **IXgen (new!)**: Indexing and marker management tools for FrameMaker.
- **User Guide Templates (new!)** (group product): For FrameMaker and WebMaker.

4. Communications and Networks

4.1 Computer Networking at CERN

Mike Gerard, IT/CS

Will *you* please help us to improve network performance?

Why should we need to ask you?

Over the last few years there has been an explosive growth in the use of the CERN General Purpose computer network and the number of systems which are connected to it. The inevitable result is that it is no longer possible to have a central support group which does everything for everyone: manpower requirements for such a service are simply not available. Thus, as with creating documents, end users have to become directly involved to some extent.

Another reason is that modern PCs and workstations have increased in power inexorably fast. In order to be able to handle this increase the CERN physical network has been modernised with the introduction of structured cabling. With this introduction every office has several network outlets, each capable of handling a connection at high speeds (100 megabits/second or above).

This said, the actual performance as seen by the office system depends not only on the user's office equipment but also upon the network equipment to which all of the network outlets are connected. This network equipment is normally a 10 megabit/second hub, on which the network access is shared between all of the connected users. Thus, most office users are being provided with a connection capacity of 10 megabits/second and shared with 20-40 other users. This is already a great improvement on the previous situation, and in most cases is sufficient for the average traffic quantity generated. There is, of course, much talk about faster networks—100 megabits/second, even a gigabit/second—and IT/CS Group is very active in obtaining and testing this new technology on the installed structured cabling.

Therefore, the situation today is that the structured cabling has given us the possibility to offer real high-speed networking to most of CERN. However, most network users are still sharing 10 megabit/second network segments with many other people and yet have the computing performance to use all of the available capacity of their shared segment. The conclusion from this is that computer users in CERN, and even outside, should help ensure a fair usage of available resources.

The rest of this article details ways in which *you* can provide such help.

A brief note on the new structured cabling

A few years ago all of the general networking at CERN was done using coaxial ethernet cables linking a chain of computer systems together. Because any errors in, or breakage of,

the chain could have disastrous effects it was decided to re-wire most of CERN with so-called structured cabling. In this scheme every office has a number of network outlets which are connected to a single star point in the same building or in a nearby one. Inside the star point there is then a patch cable which can connect outlets that are required to be operational to a particular network hub. Each hub (or pair of hubs) may offer a particular service type.

There are different types of service offered, according to the needs of the computer system concerned, and this is done by connecting the computers to multi-user hubs which are in turn connected to routers. The routers ensure that on any given type of service only the network protocols required by that service are permitted, thus eliminating unnecessary traffic.

The introduction of routers normally follows some weeks after the installation of the structured cabling. Many users will have noticed this introduction because of having been required to change the configuration of their system(s). The easy way to know if routers have been already introduced is to check if the Internet Protocol (IP) Address of the computer system is of the form 137.138.x.y, rather than the previous 128.141.v.w.

We need to know what equipment is connected to the network

Yes, we really do need to know the details of what equipment is connected onto the network, where it is, what it is, whom it belongs to and so on. This is not just silly bureaucracy, but a real need allowing us to identify any network-connected equipment which is causing problems, or even suffering because of others, and explain the problems to the named Contact Person and Main User. Without such information all that we can do is to block network access to the problem system, when it can be found, which pleases neither us nor the real owner.

You should note that *all* equipment which may be connected to the CERN network should be registered, including portable equipment and equipment not needing any IP number.

You may find it very revealing to use the Web to check on equipment in your own area. We recommend that you open the Web page at the URL <http://network.cern.ch/register> and then select the option/operation:

- View or correct the registered information about a device.

In this page you can choose to search using various criteria. As well as looking at specific equipment by name we would

strongly recommend that you look for all equipment registered for your name and for all equipment in your office. We guarantee that many people will be surprised by some of the results!

Why not do it *now*?

When you need to move equipment around!

Those of you that follow the above suggestions for checking on your network equipment may quickly realise that we have much wrong information in our equipment and cabling data bases. In part this is because it has been customary to move equipment around inside CERN without informing anyone. Unfortunately, this is no longer possible! Even in a single office, where all the network outlets may look the same, it is quite likely that moving an equipment from one outlet to another will mean that it no longer works properly, perhaps not even at all.

To permit the updating operations needed when moving equipment around we have included in the Web form at the URL <http://network.cern.ch/register> specific tools entitled:

- New Connection of a device.
- Move and Reconnection of a device.
- Removal of a device.
- Service Change of a device.
- Portable registration.

To avoid problems it is essential that these forms be used : the subsequent action, which may involve re-patching connections in a star point, will normally be done inside a few working hours.

Computer networking internal to CERN

As already stated, even with the structured cabling most users are in a situation where they are sharing an ethernet with around 20–40 other users. This is already a great improvement on the previous situation of users sharing with up to several hundred other users with very little protection against faulty hardware or software elsewhere. However, the ethernet segment is still limited to 10 megabits/second total throughput even if the ethernet interface in the user's equipment is capable of supporting the 100 megabits/second Fast Ethernet.

If this sharing and/or speed limitation becomes impossible to tolerate then IT/CS Group can advise upon ways to improve the situation, subject normally to the users requesting the improvements having a budget available to pay for such improvements.

Where no improvement is immediately possible there are a number of simple rules which, if followed, can help considerably:

- For transfers of large files, and in particular system backups which send very large amounts of disk information to a backup server, please try to schedule

these outside normal working hours. A random time between 23.00 and 07.00 is recommended (not everybody at midnight, please!).

- For systems using large remote files as if they were local ones, and especially for “diskless” systems, please be aware that every file operation can involve substantial network traffic.
- Whenever possible, if a Web browser is available on your own equipment then use that one rather than one which executes elsewhere and exports its results onto your screen. Some Web pages involving animation can cause considerable network traffic (even if the browser window is iconised!).
- Never use any screen locking device (or screen saver) which puts out continuous patterns, unless the program runs on the local system.

Access to the external world

Although network communications to the rest of the world are improving they remain both costly and offer much less capacity than the network inside CERN. This is particularly true of the communications to the USA, to where we have just a 2 megabits/second line. However, the explosive growth in Web traffic, plus the need to transfer large files of raw physics data to various USA Laboratories, means that this line can easily be saturated. In fact, the line is often saturated, as is clearly seen by looking at the the statistics on-line at URL:

<http://sunstats.cern.ch/mrtg/cernh8.html>

Again, there are a number of simple rules to help keep this USA link situation reasonable:

- For any transfer of large data files try to ensure scheduling outside peak CERN (and USA) time. Early morning (say 04.00 to 07.00) is a suggestion.
- Before going to well-known USA sites (such as *Microsoft*, *Apple*, *Netscape* or the Linux site) in search of software please check whether the required software already exists here at CERN. Even if it does not then look for available European mirror sites, since we have greater network capacity to Europe than to the USA.
- When reading news groups, either with a Web browser or with a specific news reader, do not configure it to go to a USA site. CERN already imports all reasonable news groups which may be found on USA news servers.
- Please avoid the use of Web sites offering audio and video, whether interactively or as files, if the amount of information may be large and not reasonably connected with CERN activities.

Our actions in case of observed network problems or excessive usage

Despite our best efforts, and those of the great majority of users, there are times when our overall monitoring indicates that there are either problems or some excessive use of network facilities. In such cases we will always try to contact the person(s) responsible for, or affected by, these problems or excessive usage. Our experience is that those responsible for such systems are always willing to attempt to track down the causes of the problems and eradicate them.

If you think that there are network-related problems

As an end user you may sometimes have felt that some task that you are trying to perform, and which requires network traffic, is operating much more slowly than usual. You may often feel that it is due to a network problem - some PCs with *Microsoft Windows* have been known even to say this (often wrongly!). We, as network operators, know that in many cases there is absolutely no network problem, but we also know that such a reply would not solve your problem. Therefore we are always willing to try to help identify the real problem.

You, in turn, can help us in a number of ways. Try to be quite clear on which is your own system, to which other sys-

tem is it trying to communicate and for what purpose. You can also help by checking if any other nearby people have similar problems.

The information that you gather can be collected and sent to us by electronic mail to netops@cern.ch: this is what we would prefer if the problems are not critical. If they are critical then you can call us at 74927. In the latter case remember that if lots of other people are seeing the same problems then we may be forced to leave the phone to accept recorded messages only while we deal with the problem.

If you want to know more

There is a large amount of information available via the Web pages of CS Group, specifically including more detailed information on structured cabling and the Routing and backbone project, see the URL

<http://network.cern.ch/>

So feel free to browse as much as you like (or as much as the network will let you!). In cases where this information is insufficient, or if you wish to know more about the possibilities for improvement of your own part of the network infrastructure, please feel free to contact either Network Operations (netops@cern.ch) or myself (Mike.Gerard@cern.ch).

4.2 DECnet Terminates to INFN and UK

Denise Heagerty, IT/CS

INFN terminated their DECnet access to CERN on Friday Feb 20 1998. This also terminated DECnet access to the USA and Japan.

DECnet access to the UK is expected to terminate early in March 1998.

The eventual termination of DECnet in these countries had been announced sometime ago and the Internet recommended for networking access.

Thanks to everyone who collaborated in the successful running of these services for many years.

5. Application Software & Databases

5.1 CERNLIB 98 Release

Ian McLaren, IT/ASD

As announced in the CNL 229 ("CERNLIB 98 Release", page 16, and "Cernlib on Windows NT/95", on the same page) release 98 of the CERN Program Library has taken place in March 1998. *Tar* files should be installed in `asisftp.cern.ch` at the beginning of April.

More details of the release can be found at the URL `http:`

`//wwwinfo.cern.ch/asd/cernlib/schedule.html`

This page has been updated to reflect the final system and compiler levels that will be on the supported platforms for the immediate future.

If you have any comments, please send a mail message to `HepLib.Support@cern.ch`

5.2 PAW Release Highlights

Olivier Couet, IT/ASD

The new version of PAW (2.09/17) is part of the 98 release of CERNLIB. In this version you will find several improvements and bug fixes in the various packages used by PAW. The complete release notes are available on the Web, at the URL `http://wwwinfo.cern.ch/asdcgi/`

`listpawnews.pl/paw.news20917`

The source files, binaries and libraries are available, via anonymous `ftp`, on the `asis` system:

`ftp asisftp.cern.ch`

5.3 Replacement of the CS2-PIAF Service

Dino Ferrero-Merlino, IT/ASD

COCOTIME has decided to allocate the entire CS2 computer to the NA48 experiment. Thus, the CS2-based PIAF service must be migrated to another machine by the beginning of April in time for the NA48 startup.

A new machine, a 4-processors Sun 450 system with 1 GB of main memory and 170 GB of disk space, has been acquired and will be dedicated to the PIAF service. Its characteristics are such that there should be an improved service compared to that provided by the current 8 CS2 nodes. The fact that the basic operating system and CPU family remain unchanged means that one can expect a smooth transition both from the user side and our side. The results of our tests so far are consistent with this optimistic view.

In consultation with major PIAF users the final migration date has been chosen to be April 2nd, after the main spring conferences and the LEPC.

We will try to accomplish user migration in one step: the old service will be stopped on April 1st for approximately two days, during which data will be transferred to the new system. The new system should be available for the first test users (whom we will contact directly) at the beginning of March.

Details about the migration will be given in the PIAF "message-of-the-day" as well as in our Web page: `http://wwwinfo.cern.ch/asd/paw/piaf/index.html`

6. Text Processing

6.1 New versions of T_EX and latex2html

Michel Goossens, IT/ASD

It is our intention to install the latest version of T_EX (based on Web2c 7.2) and related programs before the Summer holidays. This will bring us in line with the latest release of the T_EX Live CDROM, which will be available around Easter, thus making it easy to emulate the CERN setup in other Institutes (at CERN we only added a few CERN-specific files to the generic setup on that CDROM). A *plug and play* ver-

sion of the T_EX hierarchy at CERN is also part of the Linux CDROM produced by the Linux Interest Group.

At the same time we shall introduce the latest version (98.1) of latex2html, the well-known and much-used L^AT_EX to HTML converter. It has been much expanded with respect to the version which is currently available at CERN and should cure many of its known shortcomings.

6.2 The Web Prepares for the Future

Michel Goossens, IT/ASD

If you are a regular visitor of W3C's (*World Wide Web Consortium*) web page (URL <http://www.w3c.org>), you will have seen that in recent months many of the initiatives I mentioned in my CNL 227 article "Hyper-activity in the Web-world" (page 11–13 or available also at the URL http://wwwinfo.cern.ch/cnls/227/art_xml.html) have seen a lot of progress. In this article I want to go into some detail about the most important points.

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. Relevant URLs can be accessed via the Web Office's Web page (<http://www.cern.ch/WebOffice/>)

HTML and the downfall of the Web

The reason why HTML is so popular has much to do with its intrinsic simplicity (it is so easy to learn), but also because of the many *non-standard* extensions which are offered by the various browser vendors to help users to make their pages look *professional* and *attractive*. However, this tower of Babel of incompatible extensions is a real threat to the integrity of the Web, since it kills the universal availability of the information.

Most people love HTML because it is a clean little language that they can master in an afternoon. It is universal and should run everywhere. In the *real world* one is often confronted with broken links and a lack of portable ways to format the information. Many of us have had to (mis)use tables, frames, Java and other scripts to get a representation we like, due to the lack of a real tool to craft universally displayable Web pages.

It is probably worthwhile to look at the problem areas where we think HTML should be improved.

- *Invalid HTML*. Many commonly-used utilities produce invalid HTML or introduce vendor-specific ex-

tensions. Moreover, most users do not validate their HTML source code and browsers do not object to invalid HTML, most of the time just skipping the information which doesn't make sense syntactically. This makes it especially difficult to get consistent results between Web browsers and across platforms.

- *Broken links*. Whenever a Web page is deleted, or moves to a different host, all URL references to that page are invalidated. There has been talk about *Universal Resource Names*, which would address pages by names, and provide a level of indirection to cope with mapping names on physical addresses (just like name servers for Internet addresses), but very little progress has taken place in this area.
- *Fixed grammar*. The element and attribute set of HTML is fixed (HTML is said to have a fixed grammar, described by a *Document Type Definition*, DTD for short, the formal specification describing the syntax of a language). One thus cannot adapt the language to cope with a specific set of new applications or extend its functionality to deal with new Web technology. In the past browser vendors have added their own extensions, with the result that Web pages can only be optimised for one browser, while the others cannot display the information fully.
- *Limited support for meta-data*. There is no standard way to differentiate information, meta-data (data describing data such as keywords), and presentation inside an HTML document. Because of this lack of separation between content and form it is difficult for search engines to extract important key information from a document or to re-use the HTML source in a variety of ways, for instance, in the form of different presentations.
- *Absence of structural tags*. HTML tags are pretty good at conveying the appearance of a document, but there is

no explicit support for describing structure. This relative flatness makes it difficult to navigate through a tree (or network) of documents.

- *Data exchange difficulties.* Because of its closed tag set aimed at presenting information on the Web it is almost impossible to extract data according to tagged data fields. Moreover, just the Latin 1 character set, which only supports Western European languages reasonably well, is generally available. Thus use of HTML for other languages has to be based on extensions, preventing easy document interchange.
- *Absence of modern features.* As with any standard (even one coordinated by the Web Consortium, which responds relatively quickly to common practice) many *modern* ingredients are lacking. Amongst these are ways of refreshing information on the client side, exposing information present in dynamic entities, such as Applets, and the unavailability of an object model.

Recent work has tried to address one or more of these problems. One approach was to increase the functionality of HTML, and therefore HTML 4 was developed. To better separate form and content the style sheet language CSS (*Cascading Style Sheets*) standard was recommended. The XML (*Extensible Markup Language*) effort deals with application-specificity and better data organisation. Dynamic HTML (DHTML) goes some way towards adding a dynamic representation to Web pages. In that context DOM (*Document Object Model*) is bound to play an important role by allowing programs to access HTML (XML) elements as a structured collection of object data, each having a set of properties and methods. We shall look at some of these developments below. One must bear in mind, however, that very few browsers, if any, support these new features at present. Nevertheless it is important that users of the Web (aren't we all?) are kept informed about this evolution and have some idea of what will become available in six months to one year. This will allow us to better plan for the future and not invest unnecessarily in techniques which will be deprecated or replaced before long.

HTML 4: a richer and more coherent language

On the 18th of December 1997 W3C issued HTML 4.0 as a *W3C Recommendation*, which means that HTML documents should henceforth be marked up according to that specification (the full document is over 360 pages and is available at the URL <http://www.w3.org/TR/REC-html40> in various formats). The more significant changes with respect to the previous version 3.2, which was released in January 1997, are listed below.

- A more complete model for tables.
- The deprecation of most attributes that control presentation (such as specifying colours using elements or attributes) and their replacement by *Cascaded Style Sheets* (see the CSS Specification (URL <http://www.w3.org/pub/WWW/TR/REC-CSS1>)).

An example of a tool that could be helpful in transforming HTML documents to exploit CSS fully is *cssize* (URL <http://lara0.exp.edf.fr/glazman/CSSize/cssize.en.htm>).

- Any element can have an ID attribute, so that it can act as destination anchor of a link.

```
<H2 id="mysect">This is a uniquely
identified section heading.
< id="mypara">This is my addressable paragraph.
...
<P>As stated in a
<A HREF="mypara">paragraph</A> which was
part of a <A HREF="mysect">section</A>...
```

- Support for internationalisation eases generating documents in most of the world's languages.
- Generic objects (images, applets, other documents) can be embedded with the OBJECT element. Should a given resource not be available another can be defined, as shown in the following example, which shows a Python applet of circulating electrons in LEP, or else and MPEG movie, or else a static GIF image, and if all that fails, just writes some text.

```
<P>
<!-- First, try the Python applet -->
<OBJECT
  title="Electrons going round and round"
  classid="http://.../CirculatingElectrons.py">
<!-- Else, try the MPEG video -->
<OBJECT data="CirculatingElectrons.mpeg"
  type="application/mpeg">
<!-- Else, try the GIF image -->
<OBJECT data="CirculatingElectrons.gif"
  type="image/gif">
<!-- Else render the text -->
  Electrons circulating in the LEP tunnel.
</OBJECT>
</OBJECT>
</OBJECT>
```

The <OBJECT> element replaces (and thus deprecates) tags like <APPLET> and <IMAGE>. Below are a few examples of use.

- Reference images and view with an external application. Above we already showed how applets and images can be displayed. One can also *declare* an object, and use it afterwards.

```
<P>
<OBJECT declare
  id="electron.declaration"
  data="CirculatingElectrons.mpeg"
  type="application/mpeg">
  Electrons circulating in the LEP tunnel
</OBJECT>
...
<P>A nice <A href="#electron.declaration">
animation of electrons in LEP.</A>
```

- Launch an applet (and write a text on browsers which do not support that kind of applets).

```
<P><OBJECT codetype="application/java"
        classid="AudioItem"
        width="20" height="20">
<PARAM name="snd" value="Greetings.au">
Java applet that greets the user.
</OBJECT>
```

- Include a file (if the file is unavailable a warning message is printed).

```
...text before...
<OBJECT data="myfile.html">
Warning: The file "myfile.html" is not
available for embedding.
</OBJECT>
...text after...
```

- Advanced features include *media descriptors*, which allow the use of device-sensitive style sheets, *event attributes*, which in conjunction with scripts allow code to be executed when a given event occurs (e.g., when a document is loaded, the mouse is clicked, etc.), and the introduction of the DIV and SPAN elements, which (when used together with the ID and CLASS attributes, and style sheets) offer authors a generic mechanism for tailoring HTML to their needs and tastes.

Today, most documents on the Web are still marked up using the HTML 3.2 DTD and no browsers fully support HTML 4 as yet (Netscape 4 and MS Internet Explorer 4 do already a good job). To benefit fully of the possibilities of HTML 4 one also needs support of Cascading Style sheets (CSS Version 1), and here also browsers still do not fully conform.

Extensible Markup Languages

Even though HTML 4 is without doubt a step in the right direction if one wants to support the Web in a standard way, it is still too limited and above all too static to cope with all of the Web's many application areas (databases, search engines, optimal presentation, professional printing, data verification). A new technology had thus to be introduced to do away with the limitations related to HTML's unalterable DTD. XML (for *Extensible Markup Language*) is the answer.

After some one year and a half of work in the framework of the W3C SGML working group, the W3C issued XML 1.0 as a Recommendation on the 10th of February 1998. It is the first in a suite of standards which will revolutionise information handling on the Web.

What is XML?

XML is, by design, a subset of SGML (*Standard Generalized Markup Language*, defined as ISO standard 8879 in 1986). SGML's scope is very broad and the language rather complex (both to learn and implement). The W3C recognised this fact and decided to develop a light-weight version, XML, which does away with SGML's rarely used and more complex features. It is said sometimes that XML offers about 90% of

SGML's functionality at some 10% of its complexity, thus making sure that the *Ten commandments* of XML (its *design goals* as specified by the W3C SGML Special Interest Group when they started their activities), were fulfilled. These goals stated that XML should be straightforward to use on the Internet, allow easy processing of XML documents (i.e., XML parsers should be easy to write), and that the number of optional features should ideally be zero. Moreover, they wanted XML to be easy to learn, and XML documents straightforward to create and modify. XML makes it easy to declare the structure of a document by decomposing it into logical *elements*. All possible relations between these elements are described in a DTD (*Document Type Definition*). Several applications of XML are being defined by W3C. Amongst them are XLL (*eXtensible Linking Language*), XSL (*eXtensible Style Language*), and MML (*Mathematical Markup Language*). At the same time a coherent data model (XML-DATA) is also being discussed. It introduces an object data model which makes it possible, for instance, to express the DTD using an XML syntax.

As already mentioned, each element of an XML document has to be declared in a DTD, which provides a formal definition for the XML language instance for the document class being considered. This allows XML parsers to check the validity of document instances marked up according to that DTD, verifying, for instance, the correct nesting levels, whether all document components have been defined, etc. Note, however, that strictly speaking, the XML does not require that a DTD be present. For instance, for browsers, it could be too time-consuming for each document to download and parse a DTD and check the document against this DTD. XML applications should make sure that all documents at creation time adhere to a DTD, so that browsers can assume that they are correct. In this case XML only wants the document to be *well formed*, and it will be up to the browser to give default interpretations for undeclared elements.

The components of XML

XML is based on the concept of *documents* composed of a series of *entities* (nowadays we would probably prefer to talk of *objects*). Each entity contains one or more *elements*, and each element can be characterised by zero or more *attributes* (properties) that describe the way in which it is to be processed. The relationships between elements and the list of their possible attributes is specified in the DTD.

The beauty of XML (SGML) is that using this mechanism of defining a language with a DTD, each institute, group, company, organisation, etc., can define its own *language* for all the different kinds of document they have to handle. By being able to choose user-friendly *markup tags*, adapted to a particular application domain or cultural environment, the use of these tags will be much easier to comprehend and the markup error rate will be substantially lower than when using a more generic markup scheme. Moreover, with the help of *intelligent editors*, that will hide the markup or else guide the user by only allowing tags possible in the current context, it will be trivial to compose syntactically correct documents.

With XML the syntax for tags and entity references (a way of including foreign components) is fixed. Elements and their attributes are entered between matched pairs of angle brackets (<...>) while entity references start with an ampersand and end with a semicolon (&...;). Comments are specified between <!-- ... -->. An example is the following trivial XML document.

```
<coolxml>XML is a cool idea!</coolxml>
```

This XML document cannot, as such, be validated, since no DTD is specified. It is, however, well-formed and complete.

If we want to become a little more ambitious, we could try and define a language to compose texts for sending invitations to our friends. We could envisage something like the following.

```
<invitation>
<to>Anna, Bernard, Didier, Johanna</to>
<date>Next Friday Evening at 8 pm</date>
<where>The Web Cafe</where>
<why>My first XML baby</why>
<par>
I would like to invite you all to
celebrate the birth of Invitation,
my first XML document child.
</par>
<par>
Please do your best to come and join me next Friday
evening. And, do not forget to bring your friends.
</par>
<par>
I really look forward to see you soon!
</par>
<signature>Michel</signature>
</invitation>
```

This document is clearly marked up. All elements are delimited by start and end tags (like <date> and </date>, respectively) and they are properly nested. There also exists an outermost *root* element, which appears only here and not as contents of any other element. We say that our document is *well-formed*. Such a document is easy to parse with a computer, one of the design aims of XML. There is, however, at least one shortcoming to this document, namely that its structure is hard to guess. We have merely indicated the semantic function of a few text strings, but it is not clear what the relation between the various document components is.

To clarify the relation between the various document elements we decide to subdivide our document onto three parts: front, body, and back, corresponding to the introductory information, the message text itself, and the closing part, respectively. We also thought it would be appropriate to emphasise a few words in the text by bracketing them with <emph>...</emph> tags. A few comment lines were added as well.

```
<invitation>
<!-- ++++ Header part  ++++ -->
<front>
```

```
<to>Anna, Bernard, Didier, Johanna</to>
<date>Next Friday Evening at 8 pm</date>
<where>The Web Cafe</where>
<why>My first XML baby</why>
</front>
<!-- ++++ Main part  ++++ -->
<body>
<par>
I would like to invite you all to celebrate
the birth of <emph>Invitation</emph>, my
first XML document child.
</par>
<par>
Please do your best to come and join
me next Friday evening. And, do not
forget to bring your friends.
</par>
<par>
I <emph>really</emph> look forward to see
you soon!
</par>
</body>
<!-- +++ Closing part  ++++ -->
<back>
<signature>Michel</signature>
</back>
</invitation>
```

It is important to note that up to now we have said nothing about how this document should be rendered. The XML instance shown above only describes the information and how its various structural elements are related. How an XML application handles these data is not specified. One must define a transformation of the various elements to an output format (via a style language, such as XSL, see below) to be able to view, print, or otherwise represent or exploit the information.

Declaring document elements

In the example above we introduced a little language to allow us to mark up invitations in a convenient, clear, and easily processable way. If we want XML applications to validate documents which we are going to write according to that specification, we have to formally define our language. As explained earlier, this is done with the help of the Document Type Definition (DTD). The DTD formally defines the grammar of your little language, in other words it describes the structural relationship between the elements and their possible *attributes* as well as the external elements which we reference. Details about the syntax of the DTD can be found in the extended version of this document on the Web (URL http://wwwinfo.cern.ch/cnls/230/art_xml_full.html).

An overview of XML's other main components

In this section we shall look at two other components of the XML effort, the *Extensible Style Language* (XSL), and the *Extensible Link Language* (XLL).

XSL

As explained before, an XML application does not know, from the markup itself, how to render a document on a output device. To also standardise within the XML context the way a document should be rendered the Extensible Style Language has been defined. Historically two *style* languages existed before the XSL effort got underway. DSSSL (*Document Style Semantics and Specification Language*) has a Scheme-based formalism which allows transformations between document types and complex output specifications for preparing all kinds of output formats, including table of contents, indexes, page headers, floats, etc. CSS (*Cascading Style Sheets*) is a W3C recommendation. It targets mainly Web-based applications, which do not need the full DSSSL machinery. Recognising this fact, the XML working group is basing XSL on a subset of DSSSL (DSSSL-O) and full CSS, allowing the basic *flow objects* of both to be used. The XSL Proposal (URL <http://www.w3.org/TR/NOTE-XSL.html>) states that XSL should have the following capabilities:

- format source elements based on ancestry/descendancy, position, and uniqueness;
- create formatting constructs including generated text and graphics;
- define reusable formatting macros;
- provide for writing-direction independent stylesheets;
- have an extensible set of formatting objects.

XSL uses a declarative syntax to deal with the rendering of most tags. When needed, scripts (written in ECMAScript, a standardized version of JavaScript) can handle complex tasks. The development of XSL is still ongoing. Nevertheless, pre-releases of software interpreting XSL stylesheets are already becoming available.

XLL


XML's *Extensible Link Language* (URL <http://www.w3.org/TR/WD-xml-link>) will, of course, still support simple links as they exist in HTML for the Web today. However, building on experience gained with *HyTime* (ISO/IEC 10744, see URL <http://www.ornl.gov/sgml/wg8/docs/n1920/html/n1920.html>) and the *Text Encoding Initiative* (TEI, see URL <http://www.uic.edu/orgs/tei/>) the document linking facilities of XML are vastly improved.

In his seminal Internet paper *XML, Java, and the future of the Web* (URL <http://sunsite.unc.edu/pub/sun-info/standards/xml/why/xmlapps.htm>) Bosak explains that XML should implement and provide a standard syntax for all classic hypertext linking mechanisms, such as

- location-independent naming;
- bidirectional links;
- links that can be specified and managed outside of documents to which they apply;
- n-ary hyperlinks (e.g., rings, multiple windows);

- aggregate links (multiple sources);
- transclusion (the link target document appears to be part of the link source document);
- attributes on links (link types).

Some of the kinds of links available in XLL are shown in the following image.



Current work and applications

Below we give a list of a few of the more important initiatives and markup languages based on XML.

- BSML** Bioinformatic Sequence Markup Language.
- CDF** Channel Definition Format.
- CML** Chemical Markup Language.
- DOM** Document Object Model.
- DRP** HTTP Distribution and Replication Protocol.
- EAD** Encoded Archival Description.
- EDI** Electronic Data Interchange using XML.
- ICE** Information and Content Exchange.
- JSML** Java Speech Markup Language.
- MCF** Meta Content Framework.
- MML** Mathematical Markup Language.
- OFE** Open Financial Exchange, electronic banking and payment protocols.
- OSD** Open Software Description.
- OTP** Open Trading Protocol, for retail trade over the Internet.
- OpenTag** Markup to encode text extracted from documents of varying and arbitrary formats.
- PIGS-NG** Metadata contents description.
- RDF** Resource Description Framework.
- SMIL** Synchronised Multimedia Integration Language.
- TIM** Telecommunications Interchange Markup.
- TML** Tutorial Markup Language.
- TMX** Translation Memory eXchange.
- WAP** Wireless Application Protocol, to exchange information over narrow-band devices.
- WEBDAB** Distributed authoring.
- WIDL** Web Interface Definition Language.
- XML-Data** XML vocabulary for schemas, to define and document object classes.

7. Questions and Answers from the UCO

A more complete collection of Frequently Asked Questions can be found on our Web page at URL:

<http://wwwinfo.cern.ch/support/direct-service/#FAQs>.

General (*Miguel Marquina and the UCO Team, IT/User Support*)

Question

Is there a version of PAW for Macintoshes (OS 7.0 or later)?

Answer

This is one of the most frequently asked question on PAW! Presently, there is no version of PAW specifically made for the *Macintosh*, but it is possible to run PAW on a remote host and display the PAW graphics window on a *MacIntosh*.

This can be achieved in two ways:

1. With an X11 emulator (like *macX*) on the *MacIntosh* and the X11 version of PAW running on the host. More details can be found in the "Guide for the Usage of X Window at CERN" (see URL <http://consult.cern.ch/writeup/xusage>).
2. With a TELNET emulator using Tektronix 4014/Falco graphics emulation (workstation type 7878).

Let us take this question as an opportunity to remind all PAW users (who still do not know) of the existence of very useful Web pages containing the "PAW Frequently Asked Questions" (URL <http://wwwinfo.cern.ch/asdcgi/listpawfaqs.pl>). This page offers a keyword search facility. In particular, a search with as keyword "Macintosh" will direct you straight to the answer to the above question (see URL <http://wwwinfo.cern.ch/asdcgi/listpawfaqs.pl/5>).

Question

Regarding e-mail with pine, I find it very convenient to have my INBOX visible both from the PC and from UNIX. Can I have other folders too? At the moment I only see on the PC what I sent from the PC and on UNIX what I sent from UNIX.

Answer

Yes, it is possible to have all your folders (not only the INBOX) shared. That is, resident on the MailServer rather than on the service from which you send mail.

Let us assume that your mail loginid is "smith". Here is what you have to do on *both* sides with pine:

- Select "S (SETUP)", then "C (CONFIG)" from the main menu,
- Change the option called "folder-collections". It is probably written something like:
`<No Value Set: using "mail/[]">`
 or similar. This means access to "local" folders. You should change it to
`"{smith.mailbox.cern.ch}mail/[]"`

- Save this by accepting the changes after selecting "E (Exit Config)".

From this moment on you will use only "remote" folders, and it will tell you so when browsing through them, with the text Default for Saves(Remote). If you want to still see the "local" folders you created up to now, define folder-collections to be: `"{smith.mailbox.cern.ch}mail/[],mail/[]"` so the default will be the remote ones but you will see also the local ones.

Also notice that the collection of "sent mail" is called by default sent-mail on UNIX and savemail on PC. To make them match you might want to change the option default-fcc on one of the sides.

Question

Does there exist on RSPLUS a command similar to `lltail` on CERNSP?

N.B. When I try to run `lltail` on RSPLUS I get:

```
sh: /usr/lpp/ssp/kerberos/bin/klist: not found
Please enter your password
sh: /usr/lpp/ssp/kerberos/bin/kinit: not found
```

Answer

This difference in the Public Batch service is linked to the migration from CERNSP to RSPLUS which took place at the beginning of the week (2nd March 1998).

This was explained in the news item announcing the change (Subject: New Public Login Service to replace CERNSP: Please Test, Date: 11 Feb 98 16:12:38 GMT) in the following section:

There is one difference which may affect users submitting jobs to the Public Batch service. Although it is possible to submit jobs (with the usual LoadLeveler commands such as `lsubmit` or `xloadl`), it is not possible to use commands such as `llls` or `llpwd`.

This is also explained in Section 1.3 in this CNL.

If people really need to use these commands (*LoadLeveler* commands such as `llsubmit`, `lltail` or `xloadl`) they need to log into sp2pub (instead of cernsp or rsplus).

The administrators of RSPLUS are still doing tests to see if they can publicise an LSF command that is the equivalent of `llmore/lltail`.

NICE (Jason Oh, IT/User Support)

NICE (*NICE-95* and *NICE-NT*) provides an integrated environment for all PC users at CERN, where all applications are centrally supported, installed, and managed. It is the standard environment at CERN for *Windows 95* (*NICE-95*) and *Windows NT* (*NICE-NT*).

Below is a selection of questions that have been asked at the UCO or have arisen from the *NICE* working group. More questions (with answers) regarding *NICE95* can be obtained via URL:

<http://consult.cern.ch/uco/faqs/#nice95>

Question

How do I change my NICE-NT (Windows NT) password?

Answer

From Windows NT a feature is provided to change your password and allow synchronisation between your Novell (*NICE95* or *Windows 95*) and NT password (both passwords will be changed to the new value).

The recommended procedure is as follows:

- on logging into NT press CTRL-ALT-DEL all at once: this will invoke the security console;
- press the "Change password" button;
- key in your current NT password and select a new one, which you will have to key in twice (for verification);
- next click on "OK"; after a pause another window will appear giving the option to change your Novell (*NICE95* / *NDS tree*) password to the same value.
- **Beware:** Make sure *only* "CERN domain" is selected before pressing "OK".

This will now have changed and synchronised both passwords.

Please Note: Your newly chosen password must be a completely new password that was *never used before* on Novell or NT.

Question

I need more space on NICE/NOVELL.

Answer

You can phone the UCO Help Desk (74952): they can give you 10 MBytes extra-space (i.e., 10 MB more than your actual needs). If you need a lot more disk space (e.g., 100 MB), then you must contact your "NICE Divisional Representative". The list is given in the "If you need Help" pages at the beginning of this CNL (see also the URL <http://nicewww.cern.ch/homepage/nicerepr.htm>).

Question

How can I stop a CD from auto launching when I put it into the CD drive?

Answer

This can be temporarily prevented by holding the shift key down. To make this feature permanent from Windows 95 select from

Start -> Settings -> Control Panel -> System

Then, select device manager tab and double click on CD ROM, and do the following:

- select settings,
- and uncheck "Auto insert" notification.

From NT unfortunately the registry has to be edited: launch regedt32 and change the following value to

```
HKEY_LOCAL_MACHINE\System
    \CurrentControlSet
        \Services
            \Cdrom
```

Autorun REG_DWORD: the default value is 1, change this value to 0

Note: paradoxically Microsoft do not recommend manual editing of the registry. It can seriously "affect the health" of your PC!

Question

Is there a 32 bit FTP client available on NICE which will allow the transfer of files that exceed the 8.3 filename constraint?

Answer

Yes, the 32 bit FTP client located in

```
Start ->
    More Applications ->
        Exceed X-Terminal ->
            Exceed Applications ->
                FTP Client
```

provides this and other added functionality.

Question

Is there any documentation available regarding use of mail clients on a PC?

Answer

The "CERN Electronic Mail Guide for PC users" fully documents the IMAP clients which are available for users to connect to the CERN mail server. The recommended mail agents are *Netscape Communicator 4.04* and *PC PINE 3.96*. The Guide can be obtained from the UCO (513 R-052) or online via the URLs <http://nicewww.cern.ch/doc/maillsrve/maillsrve.htm> (English), or, <http://nicewww.cern.ch/doc/maillsrvf/maillsrvf.htm> (French).

Question

How can I “cut, copy and paste” between the UNIX and the NICE environment when using Exceed?

Answer

To copy something from the PC to the UNIX:

- Select and copy/cut whatever text you want on the PC side.
- On the *Exceed* toolbar click on “Paste Clipboard to X selection” or ...
- ... Right click on the *Exceed* icon of the Windows95 taskbar (the one which has the “Start” button) and in the menu that appears select:

```
Edit ->
  Paste to X Selection ->
    From Clipboard
```

Now, on the UNIX environment (inside *Exceed*), at the appropriate location, middle click on the mouse to paste the selection where the mouse is.

- **Note:** some mice are configured so that the middle button is simulated by pressing the right and left buttons at the same time.

To copy something from the UNIX to the PC:

- Select whatever text you want from the UNIX side (within *Exceed*).
- On *Exceed* toolbar click on “Copy X Selection to Clipboard” or ...
- ... Right click on the *Exceed* icon of the Windows95 taskbar and in the menu that appears select:

```
Edit -> Copy X Selection -> To Clipboard
```

Now, you can paste in the PC application that you want. Note that in order to paste, applications usually have an “edit” menu; if not, by right clicking on where you want to paste, you will launch an edit menu displaying the paste option.

Question

I cannot find Microsoft Internet Explorer 4 on NICE servers. This is a free product, isn't it?

Answer

Yes, this is a free product. Currently it is not certified for the NICE environment. It is strongly advised not to install IE 4.0 for there are some DLL conflicts with the NICE environment. Development is in progress to overcome this problem, but for the time being it is not officially available through NICE.

8. Computer Documentation

8.1 Book News from the UCO

Roger Woolnough and Michel Franceschi, IT/User Support

Since the last CNL we have received the following new books:

- For PCs / Windows 95:
 1. CorelDraw 6 unleashed (Coburn, Gonzalez, McCormik).
 2. Windows 95 comment faire (McFedries).
 3. Excel 97 en français, again available.
- Also for PCs / Windows NT:
 1. Windows NT4. Advanced programming (Rajagopal, Monica).
 2. Migrating to Windows NT (Heath).
 3. Windows NT User administration (Ashley, Meggitt, Ritchey).
- JAVA:
 1. JAVA 1.1 the complete reference (Naughton, Schildt).
 2. Java programming language (2nd edition) (Arnold, Gosling).
- WWW:
 1. Web Security A step by step reference (Stein).
 2. Creating Killer Web Sites (2nd edition) (Siegel).
- And some others:
 1. Red hat Linux 4.2 unleashed + CD (Pitts).
 2. C++ Object Databases Programming with ODMG (Jordan),
 3. SOFTWARE runaways (Glass), lessons learned from massive software mistakes.
 4. TCP/IP Network Administration (2nd edition) (Hunt).

5. Beginners guide to Mathematica, version 3, (Glynn, Gray).
6. A new edition of: "A Book on C" (4th edition).
7. Forthcoming arrival of the latest version of the "Java Tutorial" from the Addison Wesley Series covering 1.1/1.2, in the next weeks.

A new feature has been added to the COBS Service "Book Catalogue", which can be found at URL:

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

You can look now at the "latest acquisitions", where a list of books acquired in the last 30 days is available.

Note: If you want to propose additions to the Book Catalogue which may be considered of general interest to the user community, give feedback on the current repository or simply wish to order books, send an e-mail to COBS.Service@cern.ch.

As usual these books are available via TID, EDH or cash at the UCO Help Desk (building 513). A full list of titles, prices and stock available can be found at the above URL.

We encourage users to consult this list and check if a specific book is in "Stock" before coming along to our service (at the UCO "Help Desk").

Please also read the new instructions for book ordering which we have initiated in the New Year (1998): this is explained in detail in the article "New Book Ordering Procedure at the UCO" (page 21 of CNL229 or see the URL http://wwwinfo.cern.ch/cnls/229/art_booksell.html).

8.2 The UCO Reference Cards—Current List

Jason Oh, IT/User Support

The list is compiled regularly at URL: <http://wwwinfo.cern.ch/uco/reference-cards>

Computer Registration Documents

- 01 CERN Computer Centre User Registration (GU/SS)
- 02 CERN Computer Centre User Registration (OU)
- 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
- 50 Cross Reference for Mail Services at CERN

Text Processing

- 20 Include ps and eps files into $\LaTeX$

Text Editors

- 15 vi editor

- 16 pico editor
- 17 Emacs mini reference
- 30 Emacs and vi Commands for VM/XEDIT Users
- 34 GNU Emacs reference card

Programming Languages

- 22 SQL Basic introduction

Unix

- 08 Submitting a batch job on the SP2
- 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
- 33 Unix Commands for VM and VMS Users
- 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

- 13 UNIX connectivity from Falco terms
- 19 make command
- 35 CMZ refcard
- 36 Zephyr Messaging Service for Unix Users
- 37 Falco infinity/Setup du terminal
- 38 Falco 5220e setup du terminal
- 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

8.3 The COBS Review - Books on Java

Nicole Crémel, IT/User Support

This is the first of a series of articles where we intend to make an assessed briefing on the books offered at the "CERN COBS Service" (COMputing Books Selling Service, see URL <http://wwwinfo.cern.ch/support/cobs-service/>). This time we shall describe our offerings on Java.

Note that the list of books proposed may evolve in time according to the needs (e.g., new editions, new software or product versions, etc.) and to user feedback. Concerning Java, the book stock has recently been modified to conform to the new standard version *JDK 1.1* and all the books related to the previous version (JDK 1.0) have been suppressed.

At the date of publication of this CNL (March 1998), the list of books proposed in the category WWW. Java is (each was given a classification code: [B] for "Beginners", [A] for "Advanced", [R] for "Reference Guide", and [S] for "Specialized").

1. *Java in a Nutshell (second edition)* (by Flanagan) – [B]
2. *Java in a Nutshell - Deluxe Edition* – [B/A]
3. *The Java Programming Language (2nd edition)* (by Arnold, Gosling) – [B]
4. *Java Language Reference (2nd edition)* (by Grand) – [B/A/R]
5. *Teach yourself Java in 21 days* (by Lemay, Perkins, Morrison) – [A/R]
6. *Not Just JAVA* (by Linden) – [B]
7. *How to Program Java (2nd edition)* (by Deitel) – [A/S]
8. *JAVA 1.1 the Complete Reference* (by Naughton,

Schildt) – [A/R]

9. *Java Class Libraries (2nd edition)* (by Chan, Lee) – [A/R/S]
10. *Java Beans Developer Resource* (by Sridharan) – [A/S]
11. *Developing JAVA Beans* (by Englander) – [A/S]

1. Java in a Nutshell (2nd Edition)

(~ 600 pages / Good overview for recent Java programmers)

A quick reference with practical real-world examples and compact reference information. This book is very popular and is probably a good (and cheap) one to start with. It contains an advanced introduction to Java for C and C++ programmers, and describes the syntax of the Java language and all the classes in the Java 1.1 API (with the exception of the still-evolving Enterprise APIs.). This second edition covers Version 1.1 of the Java Development Kit (JDK).

2. Java in a Nutshell - Deluxe Edition

(~ 600 pages + CD-ROM / Complete library of documentation)

The "Deluxe Edition" of *Java in a Nutshell* brings together on CD-ROM five volumes for Java developers and programmers, linking related information across books:

1. *Java in a Nutshell, 2nd Edition* (included both on the CD-ROM and in a companion desktop edition).

2. *Exploring Java, 2nd Edition* (covers Java basics, giving a clear and systematic overview of the language).
3. *Java Language Reference, 2nd Edition*.
4. *Fundamental Classes Reference* (complete reference documentation on the core Java 1.1 classes that comprise the `java.lang`, `java.io`, `java.net`, `java.util`, `java.text`, `java.math`, `java.lang.reflect`, and `java.util.zip` packages).
5. *Java AWT Reference* (complete reference documentation on the Abstract Window Toolkit (AWT), a large collection of classes for building graphical user interfaces in Java). These last 3 volumes provide a definitive set of documentation on the Java language and the Java 1.1 core API.

This special edition constitutes a complete library for anyone doing serious programming with Java 1.1.

3. The Java Programming Language

(~ 430 pages / Good overview for recent Java programmers)

Part of "The Java Series" books from Addison-Wesley. It is a good overview of the Java programming language regarding most aspects: strategies and techniques, syntax and constructs, and the most important packages.

4. Java Language Reference, 2nd Edition

(~ 500 pages / Syntax and Language Reference - included in "Java in a Nutshell - Deluxe Edition")

This book describes all the language features and helps to understand the subtle nuances of Java, from the definition of data types to the syntax of expressions and control structures. It covers the new language features that have been added in Java 1.1, such as inner classes, class literals, and instance initializers.

5. Teach yourself Java in 21 days (Professional Reference Edition.)

(~ 1250 pages + CD-ROM / Very complete documentation for Java programmers).

One of the "bibles" for Java programming covering most of the Java topics, from the fundamental concepts of Java and object-oriented programming, to advanced Java concepts and techniques in network programming, security, native methods, and data modelling. It contains a 250-page reference section detailing the Java APIs and has been completely revised and updated to cover Java 1.1. Advanced Java topics such as development tools, the Java extension APIs, and *Java beans* are also included.

6. Not Just JAVA

(~ 300 pages / For everybody - e.g. project managers - who need to understand why Java and other Internet technologies are exploding in the software industry).

From the *SUNSOFT Press JAVA Series*. It contains a simple explanation of what is happening with Java, Internet (regarding networks) and the *World Wide Web*, and description of the main basic concepts and features.

7. How to Program Java (2nd edition)

(~ 1050 pages / Oriented towards writing applets for the Web)

A step-by-step tour of the Java programming language, more oriented towards applet programming for the Web, with lots of examples. It is designed to explain the "nuts and bolts" of the Java language, so you can make the most of your Web pages, turning them into dynamic, multimedia presentations.

8. JAVA 1.1 the Complete Reference

(~ 1000 pages / For middle- and well-advanced Java programmers of applets and stand-alone applications).

A complete reference of the Java programming language, including a comprehensive and precise description of the language and class libraries, and a complete coverage of the Abstract Window Toolkit (AWT), a large collection of classes for building graphical user interfaces in Java. It contains practical examples on all of the language aspects, as well as expert techniques for improving performance.

9. Java Class Libraries

(~ 1700 pages / Version 1.1.2 class libraries reference for both novice and experienced programmers.)

Part of "The Java Series" books from Addison-Wesley. It is a reference for the Java 1.1.2 version class libraries: `java.applet`, `java.awt` - `java.awt.datatransfer`, `java.awt.event`, `java.awt.image`, `java.awt.peer` -, `java.beans`. It contains a class index by alphabetical order.

10. Java Beans Developer Resource

(~ 350 pages + CD-ROM / For advanced Java Programmers who want to use Java Beans technology).

This book takes a hands-on approach to teaching the basics and advanced concepts of Java Beans technology with easy to understand examples and code fragments. The book is given with a CD-ROM which includes, apart from the major source code in the book (tested on Windows 95/NT and Solaris), the Java Developer's Kit 1.1.1, containing all the Beans classes and objects to build Java Beans applications (tested on Win32 and Solaris), and the Beans Developer's Kit (BDK).

11. Developing JAVA Beans

(~ 310 pages / For advanced Java Programmers who want to use Java Beans technology).

This is a complete introduction to Java's component architecture or *Java Beans*. The book discusses event adapters,

serialization, introspection, property editors, and customizers, and shows how to use Beans within *ActiveX* controls.

JavaScript and Web programming

We want to conclude this article by trying to emphasize that users should not confuse *Java* with *JavaScript*. They are two distinct products, even though they share some functionalities (i.e., programming for the Web) and have a similar syntax. We can mention the existence of two books in the UCO

catalogue in the categories WWW.HTML and WWW.Other related to this subject:

1. *Using HTML 3.2 JAVA and CGI* (Ladd,O'Donnell): expert coverage of advanced Web publishing technology, including HTML, Java, CGI, VBScript, JavaScript and more.
2. *JavaScript : The Definitive Guide. (2nd edition)* (Flanagan). This is a reference guide for *JavaScript*, and the only one presently sold at the UCO.

8.4 UCO Book Catalogue

Michel Franceschi and Roger Woolnough, IT/User Support

The **CO**mputing **B**ook **S**elling Service is located at the UCO premises in building 513 and may be used exclusively by CERN users.

Our current offering of 218 Informatics books (aussi des

livres en français) may be electronically ordered and payed via TID, EDH, or cash.

A complete, with all the relevant information, is compiled regularly at URL <http://consult.cern.ch/books>.
