


CERN COMPUTER NEWSLETTER



CERN - IT Division
CH-1211 Geneve 23
SWITZERLAND

April – June 2001

CERN-CNL-2001-002
Vol. XXXVI, issue no 2

Contents

If you need Help	i
Editorial	iv
1 Announcements	1
1.1 Free Access to Gartner Group Reports within CERN	1
2 Special 35th Anniversary	2
2.1 30 Years of Computing at CERN - Part 1	2
2.2 GISMO Project: Historical Details	6
3 Physics Computing	7
3.1 US Award for CERN Computing	7
3.2 European Union Gives 9.8 million Euros for GRID Development (January News)	8
3.3 Move from HPSS to Castor : Will the hsm Command be Transparent?	9
4 Desktop Computing	10
4.1 NICE (Windows) 2000 Migration News	10
4.2 Usage Generalisation of a <i>Problem Report Management System</i> for Problem Tracking in IT Division	11
5 Internet Services and Network	13
5.1 Computer Security at CERN: Risks are Increasing	13
5.2 Termination of LocalTalk Networking	14
6 Scientific Applications and Software Engineering	15
6.1 2001 CERNLIB Release	15
6.2 Future of CERNLIB	16
6.3 Software Development Tool Service	16
7 Desktop Publishing	17
7.1 XML Applications at CERN	17
7.2 The latest L ^A T _E X	17
8 The Learning Zone	18
8.1 Computing at CERN, Once Upon a Time...	18
8.2 Questions and Answers from the Computing Helpdesk	19
9 User Documentation	24
9.1 User Support Book Catalogue	24

Editorial Information

Editorial Board

Chief Editor Nicole Crémel (CNL.Editor@cern.ch)
Associate Editors Charles Curran
Miguel Marquina
Les Robertson
Technical advisor (printed edition) Michel Goossens

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

General Information

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

The policy is to have approximatively 4 CNL editions per year (Spring, before and after the Summer, and before Christmas shutdown).

Issue number	Deadline
Autumn edition	Friday 31st of August 2001.

Information for contributors

The CNL is processed using XML technology and templates. Articles marked up in HTML are preferred, but plain text is also acceptable. A Web form to be used for article submission (contributions) is provided at the CNL "Home" page (<http://cern.ch/ref/cnl>), link "Article Submission form" (<http://cern.ch/ref/cnl/submit>).

Help in using this form is provided via the "Help" link (<http://cern.ch/ref/cnl/help/submit>) at the top bar.

Distribution

- An electronic mail is sent to the "CNL Announcement Distribution List" (<http://cern.ch/ref/cnl/help/announcement>) as soon as a new issue of the CNL is available on the Web.
- A paper version is sent to the users who have subscribed to the "CNL Paper Version Registration" (<http://cern.ch/ref/cnl/help/register>) list.

All *subscriptions to our distribution lists* can be made through the CNL "Help&Info" page at the URL <http://cern.ch/ref/cnl/help>.

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

The compressed *PostScript* file containing the complete printable version of this CNL is available via the Web at the URL <http://cern.ch/ref/cnl-2001-002/main.ps.gz>. It is possible that, in order to optimize typesetting, the PostScript printable version differs slightly from the master HTML version. In any case the HTML version on the Web must be considered as the reference.

If you need Help

FRONT LINE Services

	Phone	Location	E-mail Address
Computing Help Desk	78888	513/R-052	Helpdesk@cern.ch
Opening hours: 8:30–17:30 (Monday–Friday)			
Satellite Help Desk	78888	52/1-004	Helpdesk@cern.ch
Opening hours: 8:30–12:30 (Monday–Friday)			
Administrative Applications (EDH, BHT, etc.)	79933	5/R-021	Ais.Support@cern.ch
Opening hours: 8.30–12:30 and 13:30–17:30			
Central Computer Operators (24/24 hours on call)	75011	513/R-049	Computer.Operations@cern.ch
General Network Operational Problems	74927		Network.Support@cern.ch

Contacts at CERN

The “Computing Help Desk” is the recommended central point for help on all computer-related issues. Following is a non-exhaustive list of individual people responsible or assigned to a given service. If you do not find a pointer to the service you need, then contact the “Computing Help Desk” for submitting your request. For all the services listed below please use the generic electronic mail address whenever possible. These mailboxes will be read even when the usual specialist is absent. It is helpful to use relevant keywords in the subject field of your message. Contact specialists directly only for very urgent cases or for very general consultations. Please note:

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

USER SUPPORT

Service Definition	Name	E-mail Address
User Relations and Service Manager	M. Marquina	user.support@cern.ch
Help Desk Manager	R. Woolnough	user.support@cern.ch
User Registration	M.C. Perler	user.registration@cern.ch
Computing BookShop	J. Megies	bookshop@cern.ch

CENTRAL SERVICES

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

See the complete and up-to-date list of contacts at URL:

<http://cern.ch/it/pdp/is/main/experiments.html#PUBSERV>

SHIFT (Reconstruction & Analysis Clusters)	T. Smith	tim.smith@cern.ch
CORE Operations	M. Vergari	mario.vergari@cern.ch
Work Group Servers	T. Smith	tim.smith@cern.ch
Distributed File Services (incl. AFS)		afs.support@cern.ch
NAP and Parallel Applications	E. McIntosh	eric.mcintosh@cern.ch
Tape Allocations and Vault Space	R. Minchin	richard.minchin@cern.ch
Stager, tapes and TMS support		tape.support@cern.ch
Central Data Recording (CDR)		cdr.support@cern.ch
CASTOR HSM software		castor.support@cern.ch
HPSS Services		Hpss.Support@cern.ch
CERN-wide Print Service	J.L. Vosdey	printer.support@cern.ch

Support for Experiments

See the complete and up-to-date list of contacts at URL :

<http://cern.ch/it/pdp/is/main/experiments.html>

DESKTOP COMPUTING

Service Definition	Name	E-mail Address
NICE (Networked PCs)	F. Hemmer	nice.support@cern.ch
PC Desktop (Repairs and sales)	C. Ball	pc.admin@cern.ch
MACintosh Support	F. Ovet	mac.support@cern.ch
MACintosh Sales	W. Hug	waltraud.hug@cern.ch
X Terminals Support		xterminal.support@cern.ch
ASIS Support		asis.support@cern.ch

Operating System Support

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

SCIENTIFIC APPLICATIONS & SOFTWARE ENGINEERING

Service Definition	Name	E-mail Address
CERN Program Library	L. Besson	heplib.support@cern.ch
PAW Support	O. Couet	paw.support@cern.ch
Anaphe and GEANT		heplib.support@cern.ch
Oracle Databases	N. Segura	oracle.support@cern.ch
Objectivity DB Service		objectivity.support@cern.ch
Software Development Tools		sdt.support@cern.ch
Adobe FrameMaker		docsys.support@cern.ch

PaRC Service (Engineering Cluster):

for each application see the "mailto's" E-mail address as given by the URL:

http://cern.ch/it/ce/support_index.html

COMPUTING FOR ENGINEERING

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

COMMUNICATIONS AND NETWORKS

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

INTERNET SERVICES

Service Definition	E-mail Address
Web support	www.support@cern.ch
Mail support	mail.support@cern.ch

In addition to this generic address, the normal support lines are listed

at the URL: <http://consult.cern.ch/service/mail/problems/contacts.html>

COMPUTER SECURITY

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

Other Newsletters

Computing for Engineers (<http://cern.ch/it/ce/newsletter/newsletter.html>)

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

Administrative Information Services (<http://ais.cern.ch/newsletter>)

Editor/Distribution: AIS.newsletter@cern.ch

Some Useful Web Pages

TvScreen (Systems Status)	http://tvscreen.cern.ch
Network Connection Request Form	http://network.cern.ch/register
Remote/Mobile Computing	http://network.cern.ch/mobile
User-Oriented Catalogue	http://cern.ch/consult
User-Oriented Service Directory (*)	http://cern.ch/services
FrameMaker Support	http://cern.ch/framemaker
Text Processing related documentation	http://cern.ch/it/asdoc/textproc.html
IT Home Page	http://cern.ch/it
User Support	http://cern.ch/it/support
NICE Home Page	http://nicewww.cern.ch/
PC Support Service Page	http://pcwww.cern.ch/
PC Farms Homepage	http://cern.ch/it/pdp/pc/
UNIX Workstation and X Terminals Support	http://cern.ch/it/pdp/ose/support.html
CASTOR Query Report	http://wwwinfo.cern.ch/htbin/pdp/castor_report.pl
Software Development Tools	http://cern.ch/sdt
Computing Support for Engineering	http://cern.ch/it/ce
Joint Controls Project (LHC Experiments)	http://itcowww.cern.ch/ITC0/jcop/
Web Services	http://cern.ch/WebServices

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

Seminars and Presentations

Designation	Usual Location	Organizer(s)	Tel.	E-mail Address
Computing Colloquia	Auditorium / bldg 500	S. Jarp, IT	74944	sverre.jarp@cern.ch
Computing Seminars	IT Auditorium 31/3-005	L. Pregernig, IT	72666	computing.seminar@cern.ch
Technical Presentations	IT Auditorium 31/3-005	S. Cannon, IT	75036	stanley.cannon@cern.ch
Informatics Tutorials	IT Auditorium 31/3-005	M. Marquina, IT		it-training.service@cern.ch

Computer Resources Allocation and Divisional Contacts

Computer Resources Allocation Committee (COCOTIME):

Chairman: H. Hoffmann/DG 60/5-006 75458 Hans.Falk.Hoffmann@cern.ch

Secretary: H. Meinhard/EP 1/1-028 76031 Helge.Meinhard@cern.ch

Divisional Representatives for Computing

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

Div. Repres.	Div. Repres.	Div. Repres.	Div. Repres.	Div. Repres.
AS	AT P. Heymans	IT A. Koppanyi	ETT M. Draper	FI C. Saitta
EST Ch. Delamare	HR M. Buttner	EP A. Bjorkebo	PS J. Riche	SL M. Bornand
ST E. Sanchez-Corral	TH E. Gianolio	TIS G.R. Stevenson		

NICE Divisional Representatives

See the URL <http://nicewww.cern.ch/homepage/nicerepr.htm>

Divisional Representatives for Networking

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

The complete and up-to-date list of contacts is given in the "Communications Systems Group" Home page, at the URL: <http://network.cern.ch>.

Editorial

Following the great success of our last CNL and the Special Edition - 35th anniversary which was greatly appreciated by our readers (as confirmed by some e-mails sent to the editor), it is not really surprising to propose a much lighter edition for the Summer.

However, you will see we have still dedicated a special chapter (Special 35th Year Anniversary) to allow the continuation of more “anniversary” articles during all of the year 2001. Not very many new contributions have been submitted for this issue, but we hope to convince more authors, via this editorial, to contribute to the next two CNL issues, and provide their historical analysis of some CERN Computing domains that have not been described yet (e.g. “Graphics Computing Evolution”, “History of Databases at CERN”, etc.)

This summer edition is then a more “traditional” one with announcements of various changes affecting our central services, like the move from HPSS to Castor, the NICE (Windows) 2000 migration, the 2001 CERNLIB release and future of CERNLIB, and the end of LocalTalk networking. Other articles are informative messages like the usage generalisation of our Problem Report Management System, and advice concerning computer security.

We wish you all a good summer holiday, and hope you enjoy reading this new CERN Computer NewsLetter edition.

Nicole Crémel, CNL Editor, IT/User Support

Hasta la vista!

As computing becomes a more and more widespread and complex activity in the laboratory, the need will increase for a means to have a wider general circulation of background information about different aspects of computing activities than is possible with the present system of Computer Notices. I therefore make no apology for introducing yet another circular which will find its way on to CERN desks. Rather I would express the hope that this newsletter will prove to be a useful source of general information on computer use and performance, programming developments and the requirements of different kinds of computer uses in the laboratory. The newsletter will be edited by the Computer manager and any comments, criticism or suggestion for topics to be included should be addressed to him.

G.R.Macleod in the Preface to CNL 1, 15 February 1966

The above is the way the DD Division Leader at the time prefaced the first issue of our Newsletter. I have felt over the past years very much on the same wavelength. I have always considered the CNL as our window to the gallery, a valid channel to explain to you, the users, what we do and why we do it. It also offers a perspective of the evolution in services and offerings.

When I became Head of User Support four years ago, I pushed for an activity-oriented structure in the CNL, rather than the former unit-oriented one. Indeed, to those of you out there it is now hopefully easier to understand what is being offered, as we try to reflect the perspective of a “customer”, rather than that of a “provider”.

But as important to explain what we do, is to help you to make better usage of our services. I introduced the “Learning Zone” chapter, with the ambition (rather unfulfilled) of offering modest training and “academic computing curiosities”. I am a firm believer that a trained Computer User is an efficient professional. Let us put behind those countless hours spent in front of the screen guessing software features by trial and error rather than by reading the documentation (when it exists and it is up-to-date!).

I also consider that history must not be forgotten nor given lesser importance just because in older times users did not enjoy a fraction of the technological breakthroughs we know today. It is precisely on such historical foundations that progress stands. The launching of the 35th Anniversary event, the historical articles on “Computing at CERN, Once Upon a Time” or the rescue of all CNL indices in electronic form may serve as samples of this conviction.

By the time you read this, I will have been reassigned to other functions away from User Support. This will likely mean that it is my last Newsletter in the Editorial Board, since my role there has been in overseeing that an appropriate reconciliation of user expectations and service provisions have been conveyed or simply expressed.

I am glad to have taken, together with Nicole Crémel and Michel Goossens, this Newsletter into the 3rd Millenium. Indeed no luck for the one who might have wanted to celebrate the next anniversary.

I’ll continue serving you though; at a different corner but with equal enthusiasm. I don’t say “Adios, but just “Hasta la Vista”.

Miguel Marquina, Head of User Support

1. Announcements

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

1.1 Free Access to Gartner Group Reports within CERN

Isabelle Forrat and Alan Silverman, IT/TS

Gartner Group is the one of the leading independent providers of research and analysis material for IT professionals. Their reports provide in-depth analysis of dominant trends, companies and products. They are an important reference resource and decision-making tool. They give you an expert look at key technologies, companies and trends in the IT industry, and they can help in the understanding of the market events and trends.

CERN has obtained a licence making these reports available online to anyone within CERN. The database contains not only current reports, updated monthly, but also some going back over a year.

The entry point is <http://cern.ch/gartnergroup/>

From here you can access information about *Gartner*, how to use this resource and the reports themselves.

Please send any feedback (positive or negative) to one of the authors of this article.

2. Special 35th Anniversary

A special chapter for authors who have "missed" the CERN-CNL-2001-001 edition to contribute some historical article to celebrate the CNL 35th birthday this year.

2.1 30 Years of Computing at CERN - Part 1

Paolo Zanella, Former Division Leader of DD

Abstract

This is the first of a three-part series, made out of the original (excellent) paper written by Paolo Zanella in 1990. I hope that it can serve future Computorians to write "IT at CERN: its birth, its life ..."

Miguel Marquina (editor)

Introduction

The origins of computing machinery can be traced back to 3000 years ago when the Chinese introduced a primitive form of abacus, or 300 years ago when Schickard, Pascal and Leibniz invented the first arithmetic machines. For us at CERN it all started 30 years ago when one of the first electronic digital computers, a huge Ferranti Mercury, was installed. It was the beginning of a success story which has changed the way to do physics and, as a matter of fact, to do any work at all in the Laboratory.

The story goes on and the changes brought in by the new information technologies continue to affect us at an astonishing rate. Those who witnessed the beginning of the information era were conscious of the potential impact of those marvellous machines on Physics as well as on Society, but they could hardly imagine what was to come. In retrospect, we can say that the 30 years which separate us from the days of those first generation machines, have been years of struggle, of sweat and tea tears, of doubt and of painful reappraisal, but also years of discovery, of excitement, of achievement and of pride to the actors in this unique adventure.

Of those 30 CERN years, I missed only the first three. I did not miss, however, the joys of programming first generation machines, although in a different environment. Actually, when I joined CERN early in 1962, the only machines around were still made out of vacuum tubes. I am, therefore, rather well placed to tell my personal account of the *what, why and how* of the evolution of automatic computing at CERN.

I shall not even attempt to be exhaustive, let alone objective. I shall tell the story as I saw it, or rather as I lived it. Instead of compressing 262800 hours into 1, I shall select and report only those events which, in my opinion, are either real milestones or make an interesting story, and I shall follow them through to their consequences, even if this means breaking the strict chronological order. The basic features of the large high-performance systems which have played a major role in the CERN Computer Center are listed in the Appendix.

The Big Bang

It took six days for God to build the Universe and a little over two years for the Ferranti engineers to produce our *Mercury*, a large assembly of complex circuitry hidden inside a row of austere cabinets, making few concessions to people's curiosity. At about the same time the PS was built in five years and today one can build accelerators 1000 times more powerful, always in about five years. The Mercury was 1000 times less powerful (16.6 kHz) than a modern personal computer which is now mass-produced by robotized assembly lines at the rate of 2 or 3 per minute!

The purchasing contract was signed on the 25th of May, 1956 and stipulated that *the machine with 1024 40-bit words of fast [120 microsec to read/write one word from the accumulator] core store and a slow [8.75 msec latency] drum store having a total capacity of 16384 words of 40 binary digits, would be shipped to CERN in early February 1957 and installation shall be completed by 1st May, 1957.... The details of the progress of manufacture will be submitted at intervals of one month.* The Mercury had a clock cycle of 60 microsec. It took 180 microsec to add and 300 microsec to multiply two long words of 40 bits. It had floating point arithmetic. Division, however, had to be programmed. It is interesting to note the absence of the word *software* from the contract, the only relevant mention being: *any programs and sub-routines which Ferranti will prepare will be made available to CERN free of charge.* Unbundling and licensing were unknown in the 50's.

The computer actually arrived during the summer of 1958 and passed the acceptance tests in October of that year. Following its successful introduction, a thick file of orders revealed the urgent need for spare valves and of *tape bins* (a critical device in those days of punched paper tape input/output!).

At the end of 1958, the availability of a programming language called Autocode, attracted the first users to the computer centre and marked the beginning of the 30 years of computing covered by this paper. It is interesting to note that the first *Mercury Autocode* compiler was written in 1956 (by R. A. Brooker), two years before the appearance of the first *FORTRAN* compiler! It had many features which appeared later in *FORTRAN*. Due to memory limitations, variable names were restricted to a single alphabetic character (5-bit Ferranti code). It is worth noting that in the mid 50's the European computer industry was still competitive. As far as competence, innovative ideas and successful products are concerned, European companies like *Ferranti, English Elec-*

tric and *Elliott* in England, *Telefunken* in Germany and *Bull* in France had little to envy their American competitors (e.g. *ERA/Remington Rand*, *NCR*, *IBM*). The size and the drive of the American market were, however, going to make the difference quite rapidly in the 60's.

The reasons behind the decision to acquire a Mercury were technical, political and financial. Technically the Mercury was definitely one of the most advanced machines around. It was considered superior to the Ferranti Pegasus, to the Elliott 404, to the English Electric "Deuce", to the Bull "Gamma 311" and to drum-type machines like the IBM 650. Ferranti was set up to produce a dozen units and CERN was going to get serial number 6. The fact that Harwell and Saclay had ordered Mercury's was given a certain weight. CERN had no experience in electronic digital computers and it was better to be part of a club of users. As to the price (one million Swiss Francs), it was five times cheaper than equivalent American machines like the ERA 1101.

Discussions went on for several months at CERN in 1955 and 1956. External experts were involved. Fifteen years after Konrad Zuse's relay-based Z3 considered by many as the first general purpose computer, ten years after the electronic ENIAC, and six years after Maurice Wilkes' EDSAC, people were still debating on the virtues of binary versus decimal machines, on the importance of floating point arithmetic, on the size of words (the Z3 had 22-bit and the EDSAC 32-bit words; the ENIAC was a decimal machine with 10 digits words. The size of central memory was 64 words for the Z3, 2n for the ENIAC and 512 for the EDSAC), and on the best way to input and output data (punched film, 5-7-8 channels papertape, or cards). Practical problems had to do with the unreliability of the hardware, the organization of the operations and the recruitment of experts. Obsolescence was already a problem.

CERN decided that the machine should be purely binary, that 32-bit words were too short, and that the speed should be as high as possible (in particular multiplication time should be under a millisecond). It also concluded that the speed and size of memory were very important (*several tens of thousands of binary digits* were a clear necessity), while the availability of floating point hardware was considered a *very desirable, although not quite indispensable, feature of a scientific computer*. But what really makes it worthwhile to have a machine, is the enthusiasm for carrying out the most difficult computations, as an American physicist put it in a letter describing his experience with a digital computer or "the courage to go ahead and solve problems which would have seemed too difficult to do otherwise."

The computers of that time had colorful names like ILLIAC at the University of Illinois, AVIDAC at Argonne, MANIAC, UNIVAC, etc... Most of them were prototypes. Everything was very much experimental. The operations arrangements at the ILLIAC were described as follows: "The user deposits his punched paper tape in a box with instructions or the operator. The operator takes these tapes, and inserts them in the machine in turn. If the code is correct, the machine delivers the answers and these may be picked up by the user the next morning. If there are errors in the code, the operator carries

out the test routines requested and the results of these tests are deposited so that these may be reviewed by the user the next morning. Thus, the whole operation of the machine becomes a fairly automatic affair." This was more or less the style adopted for our operations thirty years ago.

One of the first applications at CERN was the analysis of the papertape produced by the Instruments for the Evaluation of Photographs (IEPs), used to scan and measure bubble chamber film. The first reports convey a certain deception due to the slow tape read/write speed. Since everything had to go through the accumulator, the CPU was blocked during I/O. It was immediately clear that there was a big mismatch between the power of the computing engine and its input/output capability.

After some struggling with faulty tubes, tape bins, machine instructions and Autocode, people with lots of data discovered the existence of an IBM 704 in Paris, which offered significant advantages such as *magnetic tape units, card readers, line printers and FORTRAN!* FORTRAN II allowed 6-characters variable names and, most important, it simplified the exchange of programs with Berkeley and Brookhaven. The 1959 CERN Annual Report indicated *already* that *as the needs increase, it will be necessary to envisage the replacement of the Mercury by a more powerful system*. It was also quickly realized that these so-called electronic brains required quite a lot of human effort to be effectively exploited. Hence the proposal to double in 1960 the computer center staff (from 10 to 20).

So, by the end of the 50's, the fundamental forces, sociological and technological, characteristic of every computer service, had been discovered, including the illusion that upgrading the resources would solve all the problems and achieve the ultimate goal, i.e. *make the users happy!*

The Lessons Of The Early 60's

The next big news was the arrival of the *IBM 709*, an improved version of the 704, in January 1961. It was still a clumsy vacuum tube machine but it featured FORTRAN and all those fancy peripherals apt to improve the *quality of life*.

The word length was 36 bits, the characters became 6-bit BCD, and the core memory size jumped to 32K. The CPU was 4-5 times faster than that of the Mercury. However, to compile a *typical FORTRAN program could take several minutes!* Tape bins made way for card trays. Magnetic tape units read and wrote at 75 ips on 7 tracks and the density was 200 bpi. Peripherals were attached via their controllers to data channels. It was a significant advance in that it allowed as many as six peripheral devices to access core memory buffers while the CPU performed other work. Another important device which came with the 709 was the so-called Direct Data Connection, allowing for direct transmission of data from external equipment to memory via a channel. The speed was not ridiculous: in principle up to 1 Megabit/sec. The 709 was also equipped with one of the first *interrupt systems*.

The bad news was still the poor reliability, although the

progress was already quite substantial. Unscheduled maintenance represented 11% of the total time. Scheduled maintenance took away a time slice of similar size. So the downtime of the 709 compared not too unfavorably with the up time of the very first electronic computers.... The on-line card reader and the printer did, however, slow down the operations considerably. After one year of experience CERN added a small IBM 1401, in order to speed up the input/output, the job sequencing and the operations. The concept of SPOOLing (Simultaneous Peripheral Operation On-Line) with its 1/0 files (virtual reader/printer) has its origins in those days. Programming for the 709 was considered a difficult activity to be left to the specialists who could understand and keep up-to-date with the new techniques and the operating conventions. The machine was *an expensive resource* which had to be used efficiently. To give an idea, the list price was in the region of ten millions francs (1960 Swiss francs!). In those days a *magnetic tape cost 60 \$ (some 260 SF!)*. It was at that time that the first inescapable committees *appeared*, e.g. Computer Scheduling Committee, Computer Users Advisory Committee and the Data Handling Policy Group.

The Mercury had still its faithful users but suffered from the chronic lack of modern, fast peripherals. In 1962, as part of a lifting operation, it was enhanced by the connection of an Ampex tape unit, *compatible with IBM specifications* and operating at 3333 characters per second (over 3 times the speed of the fastest paper tape reader and some 20 times faster than a tape punch). Also, two papertape-to-card converters were installed *to ease the transfer of data from the Mercury to the 709*. By the end of 1962 it was possible to read the paper tape from IEPs into the Mercury, give it a first processing pass, write the results on magnetic tape and input it onto the IBM 709 for further analysis. The first application packages appeared at that time, e.g. THRESH and GRIND used for the geometrical reconstruction and kinematic analysis of bubble chamber events. It is amusing to note that, in spite of the growing workload and the frantic development of codes, the machines were normally switched off at weekends. But the practice of 24 hours/day, 7-days/week service was around the corner. It is also interesting to realize that things like the connection of the Ampex tape unit to the Mercury were entirely designed and implemented on site.

The next problem was how to use all those *Autocode* programs on the 709. One just wrote an *Autocode* compiler for the 709. The difficulties of developing software were soon to be learned. The first Conference recognizing the existence of a software crisis was held in Munich in 1968. Why is software always late and unreliable? People working today with modern CASE (Computer Assisted Software Engineering) or OOP (Object Oriented Programming) tools are still trying to solve the problem. But the answer in those days was: *better programming languages*.

CERN FORTRAN was defined to ensure compatibility with other laboratories and facilitate portability of codes. It was felt, however, that FORTRAN was used *mainly for historical reasons* and new, more powerful languages would be needed to fully exploit the potential of the electronic computer. As we all know, CERN was not affected by, or it missed com-

pletely, the language explosion which started in the early 60's. ALGOL, Lisp, PL/I, PASCAL, Modula II, ADA, C, PROLOG, etc... did not raise above the level of minority cultures. FORTRAN evolved through its versions II, IV, 66, 77, 8X, and it still dominates the CERN programming landscape.

It took some time to saturate the 709, but it was already clear that young physicists were becoming addicted. It was at that time that the first embryonic Operating System appeared under the name of FORTRAN Monitor System. Many other important events occurred in the early 60's, such as the connection of computers on-line to film measuring devices including the very fast automatic *flying spot digitizers* (HPD, Luciole, etc..) forerunners of the modern image digitizers and the first attempts to connect computers directly to experimental equipment (on-line experiments). The IBM 709 was operated on-line to an HPD to measure both bubble and spark chamber films. In September 1963 the 709 was replaced by a 7090, a transistorized version of the same machine, about four times more powerful.

It was at that time that the investments and the efforts started to pay off. Over 300 000 frames of spark chamber film were automatically scanned and measured in record time using an HPD Flying Spot Digitizer on-line to the 7090. At about the same time computers were connected on-line to experiments to monitor the equipment and to collect digital data from the first filmless detectors (e.g. sonic spark chambers) onto magnetic tape. The first successful demonstrations with fully auto automatic digital pattern recognition showed that computers could be programmed to replace slow human operators in a variety of tasks. Stories about computers doing things faster, better and more reliably than human beings got around producing the usual mixture of emotional reactions. In 1970 the European Physics Society held a Conference at CERN on the "Impact of Computers on Physics" and I remember the reassuring statement of an eminent physicist that "so far computers have not significantly contributed to any discovery". It was going to take another decade to see the HEP community wholeheartedly accepting the computer as a critical component of their research and admitting it to their current technological foundation trilogy: *accelerators, detectors and computers*.

1970 was also the year when the "CERN Computing and Data Handling School" was launched to educate young physicists and stimulate the sharing of computing experience between high-energy physicists and computer scientists. It turned out to be an excellent idea. The School is still alive and well, fulfilling a clear need. As to the cross-fertilization across the physics/computer science boundary, it has developed into an ideal partnership.

Actually the two disciplines have influenced each other from the very beginning. It was the physicist Bruno Rossi who built the first logic circuits which then were developed into computer hardware, and physicists have always been among the most demanding consumers of computer cycles. Enrico Fermi, when asked in the early fifties which research project would he recommend to the young Italian physicists, told them to design and build a computer. Indeed, physics research could not have become what it is without the com-

puter, and conversely, the development of the computer has been deeply influenced by the needs and vision of basic research.

The study of the fundamental properties of elementary matter involves the frontiers of human knowledge and pushes the technology to the limit of what is possible. Computer scientists have been playing with models and formalisms, architectures and languages, inventing tools and methodologies of a rather theoretical nature and they have been sometimes accused of developing general solutions in search of specific problems. When a dense problem space meets a rich solution space some good news can be expected... Evidence of synergistic effects has been accumulating ever since the beginning of the information era. *High-energy physics and information technology* are among those disciplines which, in the second half of our century, have shown the most impressive advances.

Appendix

All the major computers having served in the CERN Computer Centre in the period 1958-1968 are listed in chronological order, together with some configuration details and their characteristic features.

FERRANTI "Mercury" [1958-1965]

First generation vacuum tube machine (60 *microsec* clock cycle, 2 cycles to load or store, 3 cycles to add and 5 cycles to multiply 40 bit longwords, no hardware division) with magnetic core storage (1024 40-bit words, 120 *microsec* access time). Mercury's processor had floating point arithmetic and a B-Register (index register). Magnetic drum auxiliary storage (16 Kwords of 40 bits, 8.75 msec average latency, 64 longwords transferred per revolution). Paper tape I/O. Two Ampex magtape units added in 1962. *Autocode* compiler. At the end of its career it was connected on-line to an experiment (Missing Mass Spectrometer). In 1966 the Mercury was shipped to Poland as a gift to the Academy of Mining and Metallurgy at Cracow.

IBM 709 [1961-1963]

Vacuum tube machine (12 *microsec* clock cycle, 2 cycles to add and 15 on average to multiply 36 bit integers, hardwired division and floating point arithmetic, index registers) with core storage (32 Kwords of 36 bits, 24 *microsec* access time). Card reader (250 cpm) and card punch (100 cpm). Line printer. Magtape units (7 tracks, 75 ips, 200 bpi). Introduction of the Data Channel. FORTRAN compiler. FORTRAN Monitor System.

IBM 7090 [1963-1965]

Transistorized second-generation machine (2.18 *microsec* clock cycle) with core storage (32 Kwords of 36 bits, 4.36 *microsec* access time). Card I/O, Tape units (7 tracks, 112.5

ips, 200/556 bpi). Eight Data Channels. Interrupt System. FORTRAN compiler. Basic Monitor Operating System (IBSYS). Connected on-line to Flying Spot Digitizers (HPD and Luciole) to measure bubble and spark chamber films.

CDC 6600 [1965-1975]

Serial Number 3 (pre-production series machine). Transistor machine designed by Seymour Cray and very compact for its time. CPU clock cycle 100 *nsec*. Core memory: 128 Kwords of 60 bits. Memory access 1 *microsec*, but independent memory banks allowed for up to one access per clock cycle. Instruction prefetch. Ten overlapping functional units. Ten autonomous peripheral processor units (PPU's) each with 4K of 12-bit words core memory. Huge disks over one meter in diameter holding 500 million bits. Tape Units (half inch tape, 7 tracks, 200, 556 and 800 bpi, and one inch tape, 14 tracks, 800 bpi). High-speed card reader (1200 cpm).

First multi-programmed machine in the Computer Centre. However, SIPROS multiprogramming operating system was abandoned by Control Data. Basic SIPROS operating system had to be made at CERN. Then Chippewa OS (COS) was installed. It evolved to SCOPE which was eventually used after adapting it to CERN needs. This resulted in a non-trivial amount of changes, thus deserving the renaming to CERN SCOPE. The 6600 was connected to various FSD systems and to two on-line computers, the SDS920 and the IBM 1800, via CERN-made data links. In terms of processing capacity the 6600 was about three quarters of a CERN unit or ten times the 7090.

The change-over from the IBM 7090 was planned to take three months starting in January 1965. Major engineering overhauls had to be done instead during the first few years and ended up in a two-months shut-down in 1968 in order to modify the 6600 to incorporate logic and packaging improvements which had been introduced in the production machines. During this long period of struggling with hardware instabilities and software development and changes, computing work was done partly by sending jobs to outside computers and partly by processing data on a CDC 3400, and later on a 3800, temporarily made available at CERN by Control Data.

CDC 3800 [1966-1968]

The 3800 was a member of the 3000 series CDC family of computers, incompatible with the 6000 series machines. More conventional than the 6600, the 3800 had a 48-bit architecture. The core memory (64 Kwords) was replaced by a faster one (800 *nsec*) during its staying at CERN. This machine was eventually acquired by the State of Geneva and installed at the local University. At CERN it was replaced by a CDC 6400. It is worth noting that CERN acquired other machines of the 3000 s series, e.g. a 3100 for the FOCUS project offering semi-interactive facilities and quick sampling of experimental data at the central computers, and a 3200 for interactive graphics applications.

2.2 GISMO Project: Historical Details

David Myers, IT/CO

Abstract

This small article is triggered by the historical note written by René Brun, "The Design and Implementation of Large Packages: Hbook, GEANT3, PAW, ROOT", in the last CNL, in order to give some detail regarding the GISMO package and people who have been involved in that project.

René Brun wrote, in the last CNL:

On the software side, we had a pleasant workshop in Erice at the end of 1989. The first solid attacks against Fortran were made and the first advocates of Object Oriented Programming appeared in our field. In 1990, Toby Burnett from SLAC had an interesting prototype called GISMO of a simulation program written in Objective C and developed on a very advanced workstation (Next).

Following are some historical details (with most people involved) concerning this project.

Graphics used to mean Tektronix storage scopes until the ERASME film measurement people (Christer Ljuslin and I, based on ideas from Hans Drevermann and others) developed hardware and software for a dynamic 3D display with

a grey-level overlay. Jürgen Bettels and I then went on to develop the 3D dynamic displays used by UA1 and UA2 with Jean-Pierre Vialle. You will have seen the photos which ultimately appeared in New Scientist and Scientific American from UA1 (or maybe you are too young!).

We were joined by Leon Osterijk and started on the PIONS system, which was certainly the first Object-Oriented development at CERN. Leon wrote an Object-Oriented Fortran compiler, analogous to Objective-C, and PIONS was also a very early client-server system (like X windows) and had many other avant-garde facilities five years or more ahead of its time. (It included many ideas which later appeared in PHIGS). GKS was introduced for low-end 2-D applications, which must have been when you came on the scene, basically because we could buy software and drivers from outside.

I went on to start GISMO, with Mick Storr and Bill Atwood, later joined by Toby Burnet and many others. When this did not get enough interest at CERN I joined Wolfgang von Rueden and developed the front-end controls for the neutrino beam in C++, something which is still being used in NA48, COMPASS, HARP, and so forth. During that period much more interest developed in O-O, and projects like MOOSE were started, and later, GEANT4.

3. Physics Computing

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

3.1 US Award for CERN Computing

Press and Publications Service, CERN

Abstract

Press release issued on June 11, 2001.

On 4 June in the distinguished surroundings of Washington's National Building Museum, Les Robertson, Deputy Leader of CERN's Information Technology Division accepted a 21st Century Achievement Award from the Computerworld Honors Program (<http://www.cwheroes.org>) on behalf of the Laboratory. This prestigious award was made to CERN for its innovative application of information technology to the benefit of society, and followed the Laboratory's nomination by Lawrence Ellison, Chairman and CEO of the Oracle Corporation. Ellison nominated CERN in the Science category in recognition of "pioneering work in developing a large scale data warehouse" – an innovative computing architecture that responds precisely to the global particle physics community's needs.

CERN Director-General, Luciano Maiani said:

This is an important recognition of CERN's excellence in Information Technology. In particular, it is a reward for the teams of physicists on CERN's LEP experiments who contributed to the development and implementation of this new architecture. The prize is also an encouragement for the physicists working on the complex challenges of LHC computing.

Hans Hoffmann, CERN's Director of Scientific Computing, commented:

In addition to its major contribution to physics, CERN has been a consistent innovator in Information Technology, from the Web to its current work on Grid computing. We are delighted with this prize; particularly as it demonstrates recognition for CERN's computing initiatives not from the academic world, but from industry's leading computing experts.

The kind of computing needed to analyse particle physics data is known as High Throughput Computing, a field in which CERN has played a pioneering role for over a decade. In the early 1990s, a collaboration of computer scientists from the Laboratory, led by Les Robertson, and physicists from many of CERN's Member States, developed a computing architecture called "SHIFT"¹, which allowed multiple

tape, disk, and CPU servers to interact over high performance network protocols. SHIFT's modular design simultaneously allowed scalability and easy adoption of new technologies. Over the years, CERN has proven these features by evolving SHIFT from the systems of the 1990s, based on RISC workstations and specialised networks, to today's massive systems with thousands of Linux PC nodes linked by Gigabit Ethernet to hundreds of Terabytes of automated tape storage cached by dozens of Terabytes of caches based on commodity disk components.

CERN has since worked on evolving SHIFT in collaboration with physicists and engineers from universities and laboratories around the world. Several collaborations with industrial partners have been formed as successive technologies were integrated into the system. Today, SHIFT is in daily use by the many physics experiments that use CERN's facilities, providing a computing service for over 7000 researchers world-wide.

For the future, CERN and other particle physics institutes are working on scaling up this innovative architecture to handle tens of thousands of nodes, and incorporating computational Grid technology to link the CERN environment with other computing facilities, easing access to the colossal quantities of data that will be produced by experiments at the laboratory's forthcoming particle accelerator, the Large Hadron Collider (LHC), which will switch on in 2006.

SHIFT has proven itself as an invaluable tool for particle physics, but the award requires more than that. As Patrick J. McGovern, Chairman of the Computerworld Honors Program Chairmen's Committee points out, "Winners of the Computerworld Honors 21st Century Achievement Awards represent those organisations whose use of information technology has been especially noteworthy for the originality of its conception, the breadth of its vision, and the significance of its benefit to society". With many other fields of research such as biology and Earth observation also anticipating a growing data wave, not to mention the home entertainment industry, it is in the near future that CERN's pioneering work will bring rewards to a wider community.

The Computerworld Honors Program brings together the Chairpersons or Chief Executive Officers of the one hundred foremost information technology companies in the world to help the world's leading universities, libraries, and research

¹The beginnings of the SHIFT project were described in an article *A SHIFT in Outlook* in the previous CNL. The project began as a collaboration between the Computing and Network Division, the OPAL experiment, and Indiana University.

institutions document the global information technology revolution. The winners in each of ten categories, selected from more than 500 nominations submitted, were announced at a gala event in Washington, D.C., on June 4. Also among the winners this year was Tim Berners-Lee, who received the

Cap Gemini Ernst & Young Leadership Award for Global Integration in recognition of his pioneering work on the World Wide Web - work carried out while Berners-Lee was at CERN in the early 1990s.

3.2 European Union Gives 9.8 million Euros for GRID Development (January News)

Press and Publications Service, CERN

Abstract

Press release issued on January 10, 2001.

A 9.8 million Euros funding over three years in support of the DataGrid project was authorized by the EC Information Society Programme (within the Fifth Framework Research Programme for technology development) at the end of December 2000 and a contract has been awarded to CERN* as leader of the project.

This project is seen by international computing experts and EU authorities as an ideal test case for the development of a new model of world-wide distributed computing and the natural evolution of the World Wide Web, which was also developed at CERN. The DataGrid project was submitted to the EU on 8 May, and after a successful review the consortium was invited to negotiate a contract for 9.8 million Euros of EU funding over three years. The negotiation was rapidly and successfully concluded at the end of October 2000. The formal signature of the relative contract took place on December 29th.

As the World Wide Web is exploited by more and more people, its limitations in dealing with the huge amounts of data involved become more apparent. Its successor, the Grid, owes its name to the romance of the pioneering era when the electrical power grid was a symbol of liberty, allowing people to tap into a valuable new resource, either as suppliers or users, or both. The computing Grid shares the same pioneering spirit, transferring it to a distributed Grid of computing resources in which supercomputers, processor farms, disks, major databases, informatic systems, collaborative tools and people are linked by a high speed network.

The DataGrid project will develop and implement a novel distributed computing environment, which is specifically designed to analyse and move vast amounts of data. It will build on emerging Grid technologies, using "open source" code to create a new world-wide data and computational Grid on a scale not attempted previously, a "World Wide Grid". The resources will be made available transparently to a widespread community through layers of new "middleware", the really innovative part of the DataGrid project. Middleware could be described as software "glue", which sits between the computing operating systems and the applications, enabling collaborative working in new ways. A major activity in the DataGrid project will be the dissemination of information and experience, with a strong emphasis on ensuring that the middleware created is made available to industry, potential partners and

research areas.

The DataGrid project will provide scientists around the world with flexible access to unprecedented levels of computing resources and will initiate a new era of e-Science. It will enable next generation scientific exploration using shared databases up to a Petabyte in size (equivalent to the data contents of a pile of CD-ROMs standing about a mile high), across widely distributed scientific communities. It will allow distributed data and CPU intensive scientific computing models, drawn from the scientific disciplines of physics, biology and earth sciences, to be demonstrated on a geographically distributed Grid.

One of the first challenges for the DataGrid will be to handle the mass of data generated by CERN's next accelerator the Large Hadron Collider (LHC) which starts up in 2005. Bunches of protons will collide some 40 million times per second at the centre of huge detectors. The resulting tidal wave of data is equivalent to every person on the planet talking into 20 telephones at once. The computing power required to handle and process these data at CERN is estimated to be equivalent to about 100 000 of today's PC computers. At least three times the same power will be needed in the collaborating institutes world-wide. Clearly CERN and the four LHC collaborations cannot financially and practically sustain this effort. For this reason, CERN is placing its trust in the Grid.

The DataGrid project will help to co-ordinate national Grid projects, many of which are already underway. International connectivity will be achieved using an advanced research networking infrastructure, which will be made available by the EU Geant project. The project will explore a new scale of data-intensive Grid computing, and will provide a solid base of knowledge and experience. The middleware will be developed in collaboration with some of the leading centres in Grid technology, leveraging practice and experience from previous and current Grid activities in Europe and elsewhere.

The six main partners in the project are:

- CERN - The European Organisation for Nuclear Research near Geneva
- CNRS (France) - Le Comité National de la Recherche Scientifique
- ESRIN - The European Space Agency's Centre in Frascati (near Rome), Italy
- INFN (Italy) - Istituto Nazionale di Fisica Nucleare
- NIKHEF (The Netherlands) - The Dutch National In-

stitute for Nuclear Physics and High Energy Physics, in Amsterdam

- PPARC (United Kingdom) - Particle Physics and Astronomy Research Council

Alongside the 17 pan-European research organisations in the project, three companies are associated partners of the collaboration: CS-Systemes DSI Information, Clamart, France ũ A company which provide computer services and products for the Internet market. Datamat, Rome, Italy ũ A company which provide computer services for a large set of applica-

tions and high performance computing markets. IBM-UK, Feltham, UK - The High Performance Computing Unit of this internationally renowned computing company.

There is also an Industry Forum, which will bring together research institutions and companies from around the world with the goal to develop open Grid technologies ensuring a seamless non-proprietary Grid for all.

Further information from The DataGrid website (<http://www.eu-datagrid.org/>) or from the Project Director: Fabrizio Gagliardi/IT (Fabrizio.Gagliardi@cern.ch).

3.3 Move from HPSS to Castor : Will the `hsm` Command be Transparent?

Bernard Antoine, IT/PDP

For various reasons, CERN has decided to no longer use HPSS for storing files on tapes, and to switch to the home-grown product *Castor* (Castor fiber L.). Experiment's data already stored in HPSS is already actively being moved to Castor by the experiments themselves.

The other type of data stored in HPSS so far is the so-called "*User Tapes*" data, offered as a way to archive/restore large files (like n-tuples) on tertiary storage. This is achieved through a command called `hsm`.

Time has come ... to convert this tool to use Castor rather than HPSS, and to help users move their data.

The obvious idea would be to change the `hsm` command to be fully transparent within the conversion period:

- the `hsm put` command would start putting files into Castor instead of HPSS.
- the `hsm get` command would try and locate files first in Castor, then in HPSS, so it is the Castor copy, if it exists, which will be retrieved.

Things are already less easy with the other subcommands:

- the `hsm ls` or `hsm query` commands should look both in Castor and HPSS, and give some sort of composite output, where the only HPSS files shown are those which have no counterpart in Castor.
- the `hsm delete` command should look in both places, and delete all the existing copies of a file.

And the logic becomes over-convoluted for the command `hsm rename`. If file A exists in the HPSS tree, and file B in Castor, should `hsm rename A B` fail, even if B would have been created in the HPSS tree?

Although the design of these transparent `hsm` commands has been looked at, it is felt that the coding and debugging are not worth the required work.

For this reason, we are going to implement a simpler, translucent, rather than transparent approach:

- A copy tool is being written (it had to be anyway), and used either by the users themselves, or by the Computing Center on behalf of the users.
- While the tool is running against a specific user's files, the `hsm` command is blocked for this user, with a message asking for some patience.
- Once the copy is completed, and the files re-invigorated by a dose of Castor oil, the `hsm` command only deals with Castor files.
- This change of state will be controlled by a hidden `.hsmrc` file in the user's home directory, which file should of course not be messed with.

We think that this simpler way is less likely to create problems, and will make users more aware of the change of the underlying Hierarchical Storage System.

4. Desktop Computing

Description and news concerning the CERN desktop environment centrally provided on UNIX and PCs (NICE95, NICE-NT and NICE-2000), and providing the “general-purpose computing” environment. It consists of the hardware, software strategy, and services for the basic applications, which one expects to be generally available on desktops. This environment presents the basic computing infrastructure of the laboratory.

4.1 NICE (Windows) 2000 Migration News

Frederic Hemmer, Alberto Pace and the NICE team, IT/IS

The following articles, to explain the migration process that is beginning for Windows 2000, have already been published in the *CERN Weekly Bulletin*. A training campaign organised by IT division has started and the schedule is given at the URL: <http://cern.ch/ref/tutorials/>

The Nicest Way to Migrate your Windows computer

(French version follows)

You are probably very conservative about your computer: the fewer changes you make, the fewer problems you will have to face. But sometimes, you just have to leave your fears behind and take a plunge. And if you are a PC user at CERN, now is time to plunge into Windows 2000. However, you have nothing to fear, a team of six people have been working for over a year with more than 50 servers on this project, called NICE 2000 - preparing for the LHC era while still maintaining NICE 95/NT. After Windows 3.0, 3.1, 95/NT, this is the 4th generation of NICE desktop services, relying on an infrastructure of more than 50 servers in production.

Anyone coming from Windows 95 or NT will find a very similar interface to the one they are already familiar with. With Windows 2000, CERN users will discover a more stable and reliable working environment and will have access to all the latest applications.

More than nine thousand CERN users will be able to discover this new environment in the coming months. Some 820 computers have already been immersed in the NICE 2000 world.

An major effort will be made to facilitate migration to the new environment: More than a dozen training tutorials will be given across the lab by the NICE 2000 team, complemented by Technical Training (<http://cern.ch/training>) that will provide a set of hands-on introductory courses and Web based self-teach courses on the same subjects. In addition, the CERN bookshop in building 513 has a new catalogue with plenty of books available on Windows 2000 and all major applications. And, if anyone still feels unsure about moving to Windows 2000, there is the web page (<http://cern.ch/win>) with plenty of bilingual documentation and several recorded tutorials available online. With this and the additional documentation on paper, it should not be too difficult to make the jump into the new world.

A migration plan has been prepared in agreement with the NICE 2000 divisional representatives. Users wishing

to know the exact schedule foreseen in each division can contact their representative, who is listed in the migration task force site at the URL <http://cern.ch/w2kmtf/Membership.htm>.

Comment passer facilement à Windows 2000

En ce qui concerne votre ordinateur, vous êtes probablement très conservateur: moins vous effectuez de changements, moins vous avez de problèmes à résoudre. Mais de temps en temps, vous devriez oublier vos craintes et effectuer le grand saut. Et si vous utilisez un PC au CERN, c'est le moment de vous plonger dans Windows 2000. Pas de panique! Une équipe de six personnes travaille depuis plus d'un an avec plus de 50 serveurs sur ce projet, appelé NICE 2000, préparant l'ère LHC tout en maintenant NICE 95/NT.

Les anciens utilisateurs de Windows 95 ou NT trouveront une interface très similaire à celle dont ils ont l'habitude. Avec Windows 2000, les utilisateurs du CERN découvriront un environnement de travail plus stable et fiable et auront accès à toutes les applications les plus récentes.

Plus de 9000 utilisateurs du CERN pourront découvrir ce nouvel environnement dans les mois à venir. Quelque 820 ordinateurs ont déjà plongé dans le monde Windows 2000.

Un effort important sera effectué pour faciliter la migration vers ce nouvel environnement: plus d'une douzaine de cours pratiques sera donné dans le Laboratoire par l'équipe NICE 2000. De plus, l'Enseignement Technique (<http://cern.ch/training>) va dispenser une série de cours d'introduction pratique et des cours d'auto-enseignement sur le Web. Enfin, la librairie du bâtiment 513 met à disposition un nouveau catalogue plein de livres disponibles sur Windows 2000 et les applications principales. Et si d'aventure vous hésitez toujours à migrer vers Windows 2000, faites un tour sur la page web (<http://cern.ch/win>), qui comporte de nombreuses informations en anglais et français ainsi que plusieurs cours pratiques enregistrés et disponibles en ligne. Avec ça, plus la documentation papier, faire le saut dans le nouveau monde informatique ne devrait pas être trop difficile.

Un plan de migration a été préparé en accord avec les représentants de NICE 2000 des divisions. Les utilisateurs qui désirent connaître le programme exact prévu dans leur division peuvent contacter leur représentant, qui est répertorié sur le site du groupe spécial de migration à l'URL suivant : <http://cern.ch/w2kmtf/Membership.htm>.

4.2 Usage Generalisation of a Problem Report Management System for Problem Tracking in IT Division

Nicole Crémel, IT/User Support

Following the outsourcing of the IT Computing Helpdesk to an external company in 1998 (see related articles in CNL 232 [http://consult.cern.ch/cnls/232/art_uco.html], and *Central Computing Helpdesk - Outlines and Strategy* [<http://ref.cern.ch/CERN/CNL/2000/001/helpdesk>] in CNL-2000-001), the Helpdesk has been one of the first units (together with contracted Desktop Support staff) which required the introduction of a professional *Problem Report Management System* in order to keep track of all interventions. The tool was also needed to measure performance of the company according to the *Service Level Agreement* (or SLA) as defined by the contract. At that time CERN had decided to purchase the product which seemed to be the most appropriate and widely used (also used by the AS division), which is the *Action Request* system proposed by the *Remedy Corp.* company.

Using the *Action Request* toolkits from *Remedy*, Alan Lovell (IT/TS/IR) built in 1999 a problem-tracking application named PRMS (*Problem Report Management System*). Initially designed for the Desktop Contract, PRMS was extended in 2000 to become a general tracking system, available to all CERN groups and support lines including User Support. User Support was using (until end-2000) another *Remedy*-based application fully tailored to their needs. This application was abandoned for PRMS in order to use a common problem-tracking environment inside IT. More information on PRMS is available from the Web pages for the *Remedy* service (http://cern.ch/it-div-ts-ir/Remedy/remedy_service.html).

Some users consider such a product as an overhead, making the dialog between the support service and the user less "human". Indeed the dialog is fully defined by the system, with mail templates adapted to the different situations and problem status (e.g. "Sent to External Analyst", "Requester information", etc.). It is possible to modify and improve those mail templates (according to user feedback), but obviously a unique and "standard" configuration can only be implemented per domain. E-Mails are sent by the system (From: arsystem@sunar01.cern.ch) making them look as if a machine is replying. However an important improvement has been made and makes the system quite nice: until the case is definitively closed, it is always possible to reply to the system and to this "anonymous" address. As the problem ticket number should be inside the mail "Subject" line, the reply (and all the dialog between the user and maintainer) is recorded directly into the same case id with all associated data, and the original responsible of the case automatically informed at the same time.

Together with the advantage of tracking all user queries in a unique and central database, for later analysis and searches, another important point perceived by IT division is to promote rationalisation of policies that ensure that all user questions should be treated consistently. In that direction the

Problem Report Management System (simply called PRMS) already used by the Helpdesk and the Desktop Support contract, will be very soon used by all support lines inside IT division. For this to be achieved, a new domain *IT Services* has already been created. Progressively, categories corresponding to all the different support area in IT will be introduced. Already *Mass Storage*, including AFS, Castor and Tape support and *PDP* (for system problem on the Unix central servers) categories are in use.

Note: AFS.Support has used PRMS since the beginning of its implementation to keep track of all queries. The AFS category tree was originally set in the *Desktop Support* domain and has been recently moved, for consistency, into the new *IT Services* domain. This is an internal "cleaning", fully transparent to the users.

The *User Support* group is coordinating this migration of IT division into a standardised usage of PRMS for all problem tracking inside the division. If you have any comments regarding the usage of this tool, or suggestions for improvement, you can contact the author, or User.Support@cern.ch.

Very soon, both *Oracle.Support* and *Heplib.Support* will be using PRMS. This means that messages sent by users to those e-mail addresses will automatically be registered as new cases in PRMS, with a ticket number associated to them. This ticket number is "the reference" to be used in any further communications related to the same query (e.g. when phoning the Helpdesk to ask for more information, or when communicating by e-mail, the ticket number should always be present in the mail "Subject"). Users will automatically receive an e-mail, acknowledging reception of their query in the system, with a similar template to that written below:

Your query has been received and entered into the IT Problem Report Management System (domain ...) with the identification number:

CT000...

You should receive some feedback or a possible solution soon or you may follow the progress of this case using the URL:
<http://consult.cern.ch/problem/...>

Please do not forget to use your ticket number "..."
in any further communications related to the same query.

In PRMS ticket numbers are always in the form CT00000000NNNNN (all the leading 0s are currently imposed by the system).

Please note: The Helpdesk is using PRMS. When users submit a query directly to the Helpdesk, or any e-mail address that is mapped to them for a first filtering (e.g. *Nice.Support*, *Mail.Support* or *Printer.Support*), either the Helpdesk can provide a direct answer, and then *no* ticket number is provided, or the Helpdesk will assign the case to the appropriate category (and support line) inside

PRMS. Then the user receives the acknowledgement message *with a ticket number*. It means in practice that, *as soon as a Remedy ticket number is provided* to users, then *the query has passed from the Helpdesk*. In case users are not satisfied with the treatment of their query (transmitted to another category and support line) they can phone or contact the *User Support* group (User.Support@cern.ch) who will try to analyse where the fault has occurred, and provide an improved solution. It should be clear that the Helpdesk is not

involved any more in the problem treatment, or responsible for it, once it is transmitted to another support line.

About the author(s): *Nicole Crémel is member of the User Support group in IT, and she is sharing responsibility for the supervision of operations of the Computing Helpdesk. She is also playing a major role in the Remedy (PRMS) coordination for IT division, in collaboration with the Remedy team in TS group.*

5. Internet Services and Network

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

5.1 Computer Security at CERN: Risks are Increasing

Denise Heagerty, IT/IS (CERN Computer Security Officer)

Abstract

Computer security incidents are increasing across the Internet as a whole. Intruders target academic sites, such as CERN, to profit from their powerful systems and good Internet connectivity. This article outlines some key risks and how you can help to reduce them.

Your responsibility

Almost every week new threats and incidents are followed up by CERN's computer and security specialists. Every CERN computer user and system administrator has a responsibility for CERN's computer security, as defined in Operational Circular No 5, available at the URL <http://www.cern.ch/ComputingRules>.

Protect your passwords

The most common form of computer break-in at CERN is via accounts with passwords that become known to intruders. This may be because the password has been guessed or has been discovered by an intruder. You must choose passwords which are difficult to guess and keep them private. When you are travelling and need to access the CERN network from a remote site, you should check with your contacts there for the most secure way to connect to CERN. It is not uncommon for intruders to run software which can collect passwords, either when typed on a system or when passing in clear text (unencrypted) on a network. You should therefore take care to use a securely managed computer and where possible use software which encrypts your password.

Prevent viruses and SPAM

Viruses and unsolicited e-mail, known as SPAM, remain a constant annoyance for us all. There has been a significant increase in both of these across the Internet as a whole. Viruses are often spread by innocent users opening unexpected attachments. *Do not* open attachments which you are not expecting or that you find suspicious. E-mails warning of viruses are almost always a hoax and can be deleted, along with other unwanted SPAM mails. E-mail which requests you to forward it to many people is also usually a hoax, as are e-mails requesting money or support for a cause. Delete such e-mails without answering or forwarding them. A database of known viruses and hoaxes can be checked from the CERN security web pages at the URL <http://cern.ch/security/>.

Anti-virus software, plays an important role in reducing the impact of viruses and should be run on all CERN PCs. If you suspect your system has a problem related to a virus then contact Helpdesk@cern.ch.

The CERN mail service plays a key role in reducing unwanted viruses and SPAM entering CERN. Restrictions on attachment types which regularly contain viruses have been introduced for CERN mailing lists and these restrictions will soon be extended to the CERN mail servers. Unfortunately complete prevention of unwanted e-mail will remain impossible.

Disable non-essential applications

Another common computer break-in is the exploitation of security holes in operating systems and networked applications, with web and ftp servers being prime targets. If you administer your own computer then make sure you are aware of the applications you require and disable all others. You may have applications running which you do not need. All applications have a risk of containing security holes which if exploited can impact more than your own system since an intruder can then use it to attack others. If you do not need to run web, ftp or other applications on your system then you will reduce the security risk for all of us by disabling them.

Regularly install security patches

Insecure systems on the CERN network increase the risk for the whole site. System administrators need to keep their systems regularly updated with security patches, both for the operating systems and every application service accessible on the network. System administrators running CERN-certified operating systems can profit from IT Division's services which apply tested patches automatically or provide them for manual installation. To help keep common applications secure, we have recently started network security scanning. This is initially targeted to help web server security but will be extended to ftp and other applications. Further information on network security scanning at CERN is available at the URL <http://cern.ch/security/scans>.

Summary

Increasing computer security problems are a sign of the times and we must all adapt to reduce them. Protecting passwords, not opening suspicious or unintended e-mail attachments,

preventing unnecessary access to computers and applications, and regularly installing security patches are a good

start in the right direction.

About the author(s): *Denise Heagerty is the CERN Computer Security Officer.*

5.2 Termination of LocalTalk Networking

Mike Gerard, IT/CS

Abstract

Announcement of the end of support for the interconnection of Apple LocalTalk networks, of old Macintosh computers and Apple LaserWriters, on 1 July 2001 when the interconnecting FastPath routers has been switched off.

As announced in two Computer Newsletter Bulletins last year, on 1 July 2001 we have removed the remaining *Fast-Path* routers, which interconnect the standard CERN networking with the *LocalTalk* networks which allow slow-speed networking of old Apple Macintosh computers and LaserWriters.

All owners of registered Macintosh computers still (possibly) on *LocalTalk* networks should have received a warning by e-mail at least once during the previous months. In addition, we have tried to identify all LaserWriters still operating on *LocalTalk*, so that the Printer Support service of IT Division can suggest a suitable upgrade/replacement.

Any persons for whom this action may cause problems should contact the author (Mike.Gerard@cern.ch).

About the author(s): *Mike Gerard is known as Network troubleshooter!*

6. Scientific Applications and Software Engineering

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

6.1 2001 CERNLIB Release

Ian McLaren, IT/API

Version 2001 of the CERN Program Library for the CERN supported platforms and compilers (<http://cern.ch/cernlib/news/cernlib.platforms2001.html>) was released on *Monday, 11 June 2001*. To fix the bug in PAW for some uses of the “opera/add” function an update was made on *Monday, 18 June 2001*. The basic code was frozen on 7 May 2001.

Major changes

- *Lapack* has been separated from *mathlib* and requires *liblapack3.a* and *libblas.a*, mainly to avoid name clashes with *kernlib* and *naglib*, but also to give users more flexibility in managing these libraries. For example, users only referencing *Lapack* should change from `cernlib mathlib` to `cernlib lapack`.
- *Pawlib* also references *Lapack* so users of *pawlib* applications may need to change their link sequence, the most likely being interactive *geant321* and *geant++* applications. However, the `cernlib` command has been changed to help with backwards compatibility.

Monte Carlo libraries

- There are new versions of the Monte Carlo libraries *pdflib804*, *isajet751* and *pythia6152*. As usual the previous versions can be obtained from the *obsolete/lib* directory for the relevant platform. Note that the previous *pdflib* had no version number.
- *Pythia* had dummy routines (*pdfset*, *structm* and *structp*) which prevent easy use of *pdflib804* on most Unix systems. These have been removed and made available in the new library *libpythiad.a* for those users who cannot use the recommended link sequence `cernlib pythia6152 pdflib804`.
- *Jetset74* had dummy routines (*pdfset* and *structm*) which prevent easy use of *pdflib804* on most Unix versions. These have been removed and made available in the new library *libjetsetd.a* for those users who cannot use the recommended link sequence `cernlib jetset74 pdflib804`.

Notes on Monte Carlo releases

- Some years ago, when the new version of *Pythia* was released, it was firmly pointed out by the Monte Carlo

authors and physicists from Opal and Aleph that the old Cernlib policy on these libraries was inherently dangerous and that users should be made aware by the version number which version (and Physics) they were getting. This has been applied to *Jetset*, *Pythia*, *Isajet*, *Photos* and now *pdflib*, and the user friendly but dangerous links were broken.

- So our policy is to find a compromise between the latest stable version from the authors and what is required by the CERN experiments. The previous versions are made available by keeping them in the “obsolete/lib” directory for the relevant platform.
- In the case of *pdflib*, the author requested that *pdflib804* should replace the very old *pdflib* which was created before the new policy was introduced. So, users are strongly recommended to change to *pdflib804*. If, for some reason, they can not do so, then version 7 can still be used from the old or obsolete directories.

Further details

Details of this release are at the Cernlib Web page <http://cern.ch/cernlib/>

The the 2001 release notes are at the URL <http://cern.ch/cernlib/news/cernlib.news2001.html>.

Some further interesting links are:

- List of changed routines since the 2000 release: <http://cern.ch/cernlib/news/cernlib.changelog.2000>
- PAW Release notes: <http://wwwinfo.cern.ch/asdcgi/listpawnews.pl/paw.news21222>
- List of changes to PAW: <http://wwwinfo.cern.ch/asdcgi/listpawbugs.pl>
- Platform-specific *README* files in the Installation: <http://cern.ch/cernlib/version.html>.

6.2 Future of CERNLIB

Jürgen Knobloch, IT/API

In the FOCUS meeting of December 2000 the future of CERNLIB was agreed as follows:

For FORTRAN software (CERNLIB), no new functionality and no ports to new platforms are planned. The final release in spring 2002 will be only on Linux, Solaris, HP and Tru64 UNIX, with maintenance for one more year after that. Limited support will be available until end 2003.

A special agreement for the LEP experiments was presented at the FOCUS meeting of May 2001:

- Support CERNLIB until end 2003 as previously agreed by FOCUS.
- Keep *as is* CERNLIB (or an agreed subset), ZEBRA, GEANT 3.21 on Linux *for ever* but with no new code, no user support and no port to IA64. Recompilations will be done for major new Linux versions, but must

be tested by the experiments. If the tests fail, the last working frozen version will be kept: this would be the start of a *museum* system.

- PAW support may be discontinued after 2003 if an appropriate replacement product exists.
- The CERNLIB build procedure will be further streamlined and documented for the last release in 2002, to allow non-experts to continue building of CERNLIB on Linux.

The complete minutes of the two meetings are at the FOCUS Web pages at the URLs:

<http://cern.web.cern.ch/CERN/meetings/FOCUS/Focus20/minutes20.htm>

and

<http://cern.web.cern.ch/CERN/meetings/FOCUS/Focus21/minutes21.htm>

6.3 Software Development Tool Service

Eric Poinsignon, IT/API

Abstract

Description, content and news from the service.

You can find the complete information on what we have and how to use it at URL <http://sdt.cern.ch/> and you can contact us at e-mail address: Sdt.Support@cern.ch.

Last news

- *FrameMaker 6.0*
(<http://sdt.cern.ch/FrameMaker/>)
for all platforms: (Nice 95, NT, 2000, Solaris, HP/UX, AIX, Linux via Solaris, and Macintosh) is now available. Users on Nice 95 & NT will have to install it locally via a new Nice Menu, but version 5.5.6 will stay available. Users on Macintosh will have to borrow the CDROM. SGI/Irix is not supported anymore

by Adobe. Framemaker on IBM/Aix and HP/HP/UX: does not contain “distiller” anymore, it is no more possible to generate acrobat/pdf file.

- *Together 5.0 build 1340*
(<http://sdt.cern.ch/Together/>)
is only available for beta testers as we do not have definitive licenses yet. It's a UML Integrated Development Environment generating Java and C++.
- *Rational Unify Process (RUP) 2001.03*
(<http://sdt.cern.ch/RUP/>)
is available. RUP is a Web-enabled software engineering process for software best practices.

About the author(s): *Eric Poinsignon is looking after the SDT Service run by IT/API.*

7. Desktop Publishing

Everything related to “Text Processing”: HTML issues, $\text{\LaTeX}$, Word, Framemaker, etc.

7.1 XML Applications at CERN

Michel Goossens, IT/API

In the AFS directory `/afs/cern.ch/sw/XML/cdrom` you can find an up-to-date set of publicly available XML tools for Linux as well as a lot of documentation and tutorials. The file `index.html` in that directory provides an entry point to much of what is available.

These XML tools can be executed by including the directory `/afs/cern.ch/sw/XML/cdrom/bin/i386/` in your PATH variable.

Some important tools in that directory are:

<code>batik.sh</code>	Apache's SVG browser
<code>fop.sh</code>	Apache's XSL-FO to PDF converter
<code>oraxml.sh</code>	Oracle's XML parser
<code>oraxsl.sh</code>	Oracle's XSL parser
<code>rxp</code>	Fast XML parser
<code>saxon.sh</code>	Mike Kay's XSL parser
<code>xalan.sh</code>	Apache's XSL parser
<code>xerces.sh</code>	Apache's XML parser
<code>xp.sh</code>	James Clark's XML parser
<code>xt.sh</code>	James Clark's XSL parser

To run these scripts you must set the shell variable XMLCDROM as follows (depending on your Unix shell):

```
# for Bourne-like shells
XMLCDROM=/afs/cern.ch/sw/XML/cdrom
export XMLCDROM

# for C-like shells
setenv XMLCDROM /afs/cern.ch/sw/XML/cdrom
```

The whole XMLCDROM directory tree, including binaries for Linux only, just fits onto one CDROM. A new CDROM edition is produced every few weeks and a copy can be bought “at cost” from the Computer Bookshop. A CDROM is convenient for running these tools locally (e.g. on a portable computer running Linux not connected to AFS).

The directory `win32` contains a few files that can be installed on Windows 2000 (e.g., Microsoft's excellent XSL parser `msxsl`, version 4.3.2 of Amaya, etc.). I plan to install these tools also in a Windows 2000 directory on Nice 2000. If there is interest we could also produce a Windows 2000 CDROM.

However, as the binaries take up a rather large fraction of the space, I am looking into producing a DVD containing all the binaries for Linux, Windows 2000, as well as for Solaris. Given the six-fold increase in available space it would also allow me to add other interesting utilities that at present do not fit on a single CDROM.

7.2 The latest $\text{\LaTeX}$

Michel Goossens, IT/API

The latest official version of $\text{\TeX}$ at CERN was installed in June 1999, i.e., more than two years ago.

Recently I made available an alpha release of TeXlive 6 (June 2001) on AFS, so that more recent versions of $\text{\LaTeX}$ and its packages can be used.

To have access to this release, just add the directory `/afs/cern.ch/sw/TeXlive/6alpha/bin/i386-linux` for Linux, or `/afs/cern.ch/sw/TeXlive/6alpha/bin/sparc-solaris2.7` for Solaris to your PATH variable.

One interesting characteristic of this release is that $\text{\LaTeX}$ is compiled in its “huge” variant and comes with the hyphenation patterns of over twenty languages preloaded, so that, after choosing the relevant language (e.g., with the `babel` system) you can typeset text in that language.

```
> PATH=/afs/cern.ch/sw/TeXlive/6alpha/bin/i386-linux:$PATH
```

```
> export PATH
> latex
This is TeX, Version 3.14159 (Web2C 7.3.3.1)
**\ relax
LaTeX2e <2000/06/01>
Loading CZ hyphenation patterns: Pavel Sevecek, v3, 1995
Loading SK hyphenation patterns: Jana Chlebkova, 1992
Babel <v3.7h> and hyphenation patterns for english, dumylang,
nohyphenation, czech, slovak, german, ngerman, danish, spanish,
catalan, finnish, french, ukenglish, greek, croatian, hungarian,
italian, latin, mongolian, dutch, norwegian, polish, portuguese,
russian, ukrainian, serbocroat, swedish, loaded.
```

Once this release consolidates it will be made available on Linux, Solaris, and Windows 2000 in a more standard way. The other Unix platforms being “frozen” they will *not be upgraded* and will stay with the release of June 1999.

About the author(s): *Michel Goossens is a CERN authority on $\text{\LaTeX}$, XML and Electronic Document Publishing techniques in general. He has written several articles and books on the subject.*

8. The Learning Zone

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

8.1 Computing at CERN, Once Upon a Time...

Miguel Marquina (Editor), IT/User Support

Abstract

A retrospective look at events in the CERN Computing arena.

Five years ago...

Closure of CERN Support for SunOS 4 and ULTRIX Operating Systems and SGI Desktop Systems

(CNL 223, “Migration Special - VM Shutdown”, April-June 1996)

At the present time, the UNIX Workstation Support team in CN/DCI Group offers support on a total of 7 UNIX architectures plus 2 flavours of X terminals. With ever-increasing demands on our services (the number of UNIX workstations has increased by some 50% over the past 2 years and the number of X terminals has at least doubled) and no prospect of increasing our resources, we are forced to look closely at limiting services...

Our present efforts to rationalise platforms is not really new, is it? Thank Heaven Linux is there and there is light at the end of the tunnel. I mean RedHat Linux, not SUSE, or Caldera, or... I really mean RedHat 6.1. Or rather 7 by September? Surely Sun Solaris is a very acceptable exception to it. Sorry, I mean Solaris 2.6, not really 2.7 yet. But wait a minute, we cannot afford not plunging in parallel into it, can we?

Rumours say Michael Ende might have found in this inspiration for his book... (Ed.)

Cabling, Routing and Webbing

(CNL 223, “Migration Special - VM Shutdown”, April-June 1996)

The structured cabling project is advancing well, with more than 5000 outlets installed and cabled, and more than 1800 users already connected. At the same time, the conversion to a routed network has made a good start, with the entire Computer Centre and several buildings converted so far.

And today routed network is ubiquitous on-site. Another example of the foresight of the Computer Networks staff in offering us appropriate infrastructure to the information exchange needs of the lab for the years to come.

A bow to Brian Carpenter from here my humble corner (Ed.)

Listbox for Novices

(CNL 223, “Migration Special - VM Shutdown”, April-June 1996)

There are several typical mistakes that we are seeing regularly from users of the Listbox mailing list service. The first recommendation is to read the documentation!

This and another four articles addressed topics on the then newborn Listbox service. It is interesting to see that the typical patterns we see in supporting users stand time just as solidly as rock.

Just for the anecdote, did you know that I am the father of the service (built to replace the VM's LISTSERV), including its christening? Now you do. (Ed.)

Ten years ago...

Format of Program Library source Code

(CNL 202, April-June 1991)

Up to now the public source files of the Program Library have been kept on binary PAM files... The Program Library Office is planning to make available and distribute only CARD PAM files in the standard areas.

I managed to convince Julius Zoll The Legend to drop support for the “proprietary” binary format. And how proud I was... (Ed.)

Evolution of handling Unix files with ZEBRA

(CNL 202, April-June 1991)

The KERNLIB package Z310 CFIO, to read/write UNIX files by interfacing to the C library was released [...] in March 1991. This worked correctly for disk files, but not for on-line tape files.

Since SHIFT is now refreshed in our minds (see “Prestigious US Award for CERN Computing on page 7), guess who created the ancestor of RFIO. Julius again! (Ed.)

Twenty years ago...

Coming Changes to the IBM System

(CNL 158, May-July 1981)

Last autumn, the Scientific Policy Committee agreed that it was necessary for us to add a significant amount of capacity to the IBM system if we are going to be able to meet the increasing demands coming from the start of p-pbar experiments, increased intensity at SPS and from an expansion in the scale and range of terminal services.

[...]

3081 is a new processor from IBM that uses a new technology... It will have 16 Mbytes of memory 16 channels and two CPU's that together give a capacity of approximatively. 3.8 times that of the 370/168.

[...]

At the same time as the 3081 is installed the 3032 will be dismantled. This work will almost certainly take place over the weekend of Jeune Genevois, September 10-13.

Two full mainframes relocated and practically operational in four days. CPU capacity blossoming to address computing requirements fully in the house. Those were different and challenging days.

Today we trash a PC several factors faster and RAM-richer when a farm gets an itch. But we have other challenges equally stimulating. Find the ways to make LHC Computing a truly international Collaboration because we can no longer afford doing it alone. Experimental Physics keeps teaching us the ways... (Ed.)

Thirty years ago and beyond...

FIFTY - Editorial

(CNL 50 - 16 February 1970)

The first edition (may it become historic) opened with the

statement of intention: [see Macleod's quote in the Editorial]

We hope the Newsletter in its first fifty editions has in some small way succeeded in this. Our attitude has been that the Newsletter should be just a little more than another boring trial, of the terminology that beggars any computer document. We regret that we have not found the time to present more information about the other computing activities of CERN. After all our department has fewer computers than either of the physics departments. Perhaps our regret is really that we have not succeeded in stirring up enough feeling about computers - for they dominate most people's lives and experiments at CERN. Letters have been few. We wish there had been more.

By somebody who knew what Public Relations is about. (Ed.)

8.2 Questions and Answers from the Computing Helpdesk

Collected by the IT/User Support Team

Nicole Crémel (Editor), IT/User Support

Abstract

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

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

Windows Support

(Hannes Schwarzbauer, IT/User Support)

Question [424] – Word option "strikethrough"

I want to mark some text in Word as deleted, i.e. crossed out with a horizontal bar (like one would do it on a draft paper), but still want to leave it visible.

The online help of Word doesn't seem to have any reference to it, selecting keywords like "cross out", "cancel", "overwrite", "delete",...

Answer

The online help is there, but the keyword to look for is "strikethrough".

Question [426] – Excel 97 error ("document not saved")

I'm trying to set up an Excel 97 macro, but when I want to save it the system replies "document not saved" (no other error message).

As far as I can see, the macro and the Excel worksheet look correct, without any virus.

Answer

The process was interrupted while creating the temporary file, either because you pressed the ESC key or because there was some hardware or software failure or other media problem.

The original file should be still intact and unless your computer or workstation failed, the version of the file containing your current changes is still open in memory. Save the file to an alternate drive. If your computer or workstation did fail, you can use the original file.

(Note that any changes made in the last revision will be lost.)

Question [448] – Ethernet cards on W2000

Which PCMCIA cards are allowed for a W2000 installation?

Answer

The only Ethernet cards that are currently (Dec. 2000) supported by the Win2000 installation at CERN are:

Intel EtherExpress PCI PR010+, PR0100B
3c589 3Com EtherLink III adapter
3c90X 3Com EtherLink PCI adapter

Question [468] – W2000 portable DHCP connection to CERN domain

I have connected my portable PC (running Windows 2000) to a special CERN DHCP socket foreseen for portables, but I cannot access the CERN domain or see any of its applications.

Answer

Computers need to be registered correctly in LANDB for Windows 2000 to work and deploy applications. In this case, the computer had not been registered as portable from the network database.

You can check your computer configuration going to the page at the URL <http://cern.ch/winservices> and then select "Computer account status", which will display in red incorrect settings.

The PC has to be registered correctly and you will have to wait for the active directory synchronisation to complete.

Question [478] – PC name

How can I find out the name of my PC? (It is not written on the PC).

Answer

Start - settings - control panel and double-click on "network". The "identification" frame contains the PC's name.

Note: it is good practice to stick a label with the PC name on top of the screen (for quick identification).

Question [482] – Virus check on PCs

I have received a file on my PC by mail, and I was told to first run a virus check on it before opening. How?

Answer

- click on the file;
- save it to disk;
- right-click on this saved file;
- choose "scan for viruses";
- wait a few seconds until you get a result.

If no viruses found, you can open the file by double-clicking.

Question [502] – Strange Word template layout

When looking at a given template with WORD 97 I see things such as FORMDROPDOWN and FORMTEXT and TIME \@ "dd.MM.yyyy" instead of the selection box, text and today's date, as usual.

N.B. This is independent of the machine on which I work.

Answer

This means that the document is form protected. You have to select:

- From inside Word: Tools - unprotect document
- Tools - Options - View, un-tick the check box "Field Code".

Question [518] – No standby mode for portable PC

I have an HP Omnibook 6000, running W2000. When I try to go into "standby" mode I get a message saying that

"the SMS Virtual Keyboard driver is preventing the machine from entering standby mode. Please close all applications and try again. If problem persists you may need to update the driver".

I closed all applications, but nothing changed.

Answer

This usually means the user has installed some additional software. Un-installing it should solve the problem.

In one particular case, *FrameMaker* was installed along with *Adobe type manager 4.0*. By removing *Adobe type manager* the problem was solved.

UNIX Support (General)**Question [416] – Kill process on Unix**

I cannot kill a process on ATLAS cluster. Why?

Answer

This is a known problem with AFS (November 2000) which has some difficulties to handle very large amount of files. This often results in make processes which are un-killable and the only solution is for the system manager to reboot the node where this process is running.

Question [422] – Delete Unix files starting with a dot

I want to empty completely one of my directories but rm leaves files starting with a dot, e.g. .xyz.

Answer

rm * or rm *.* will only delete the ones you can see with ls, i.e. not including the .xyz files. rm .* will do it.

Question [467] – DECserver connection (VT100/VT200 terminals)

I have a VT100 (VT200) terminal and I used it to connect to HPPLUS or RSPLUS, probably via a DECserver. It does not work anymore since January 2001.

Answer

Indeed with the VXCERN shutdown (Christmas 2000) those terminals are **not** supported anymore. The only solution to connect to CERN computing services is to change device (e.g. PC connected to the CERN Network).

Question [470] – CCDB tools - xuserinfo - xuserreg - xspaceadm

I cannot run anymore xuserinfo / xuserreg / xspaceadm on HPPLUS.

Answer

As from Tuesday, Feb.13 2001, the interfaces xuserinfo and, for the Group administrators, xuserreg and xspaceadm, have been disabled.

New versions have been introduced on Tuesday, Feb.6, 2001 on the Linux based service Lxplus; they will still be available on Rsplus to ease the transition.

N.B. The interfaces on Lxplus are nearly *identical* to the ones currently in use: the only minor differences would come from a change in the version of the Oracle tool used.

Access to Web version:

You can also use the Web version of the interfaces: to do this, you must first register yourself by means of the “Web Reg.” button which is displayed on the first page of `xuserinfo/xuserreg/xspaceadm`. Once registered, you can use the URLs `http://userinfo.cern.ch` (resp. `http://userreg.cern.ch` and `http://spaceadm.cern.ch`).

UNIX Support (Batch)

Question [420] – LSF and SSUSP state

It seems that all the batch jobs I have submitted have been suspended because a user submitted another job with some SSUSP signal.

Answer

In most LSF clusters (e.g. in Atlas), batch queues have different scheduling strategies. The policy is that shorter jobs have higher priorities than others. In practice that means that longer jobs on a particular node might be suspended (then resumed) to have shorter requests get executed. This is automatically done by the LSF system itself. That's why from time to time you may see requests in the SSUSP state like right now:

```
> bjobs -u all -m shiftatlas | grep atlas032b
556674 jleveque RUN    1nh atlas0100at atlas032b ...
556620 mommsen RUN    1nd atlas0050at atlas032b ...
556624 mommsen SSUSP 1nd atlas0050at atlas032b ...
```

Question [449] – LSF output

How I can check the output files of a job submitted with LSF while it is running?

Answer

The command is `bpeek/`.

If you look at the page at the URL `http://consult.cern.ch/writeups/plusguide/` you will find a very informative chapter on LSF. Moreover, the URL `http://consult.cern.ch/cnls/230/art_pbatch.html` has an example and also links to where you can find all the information on lsf and the online manuals.

Question [503] – Options in command `bsub`

I want to run a job on LXBATCH which requires the machine to have 200 MB of memory. Some machines in LXBATCH have only 128 MB. Can I specify the machine name (e.g. `lxbatch026`) in the command `bsub`?

Answer

In order to request a certain amount of memory for a batch job you should use the `-R` option for this special resource, e.g.:

```
bsub -R "mem > 200" ...
```

requests for the resource of at least 200 MB of memory.

Please do request a resource and **not** a host. To check, if there are hosts available for a requested resource, you can use the command:

```
lsload -R "mem > 200 && public" lxbatch
```

which shows some machines in the cluster LXBATCH, which are available for public use and which have more than 200 MB of memory.

More information on the batch system can be found at the URL `http://wwwinfo.cern.ch/pdp/lsf/index.html` and in the guides specified there.

Mail Support

Question [403] – Limited mail size for mailing lists

When I sent a message to a mailing list it was rejected, because apparently the message was too long. Why? Is it a problem with disk space?

Answer

A size limit is imposed on messages sent to mailing lists. You can get this limit when you search this mailing list with SIMBA.

SIMBA is the Web interface that allows you to manipulate your own list(s).

Generally messages sent to mailing lists should not contain attachments (they should contain just URLs pointing to the files on the web).

Question [456] – “pine” and Mailserver password

How can I use “pine” from my new workstation without needing to type my Mailserver password?

Answer

Users with personal workstations need to use `userinfo` on the web or `xuserinfo` on Unix, and set the “hosts” section. The mail team will make the necessary update to allow Mailserver access without password typing.

Note: For the central hosts the system managers update a list by hand and the mail team takes it into account.

If you see any problem with `userinfo`, i.e., if this doesn't result in a resolution by the next day, or if on a central server you are asked for a password, then contact the Helpdesk (`helpdesk@cern.ch`), mentioning the name of the new workstation added to the cluster which requires the typing of MailServer password every time you wish to use pine, or the cluster node. In that case we must ask the system managers to make the necessary manual entry in the table of cluster workstation members.

Question [475] – Outlook and attachments

If a mail client on W95/Nice (and maybe other systems) has set the default character set to “user defined”, Outlook/W2000 gets confused, shows the entire header, but not the sender's name in the list and may lose in-line attachments.

Answer

The sender must change the view/character set/ to Western and set “set default character set” to make this change effective.

Question [483] – Mail service e-mail addresses

What are the different e-mail addresses (contact points) related to the CERN Mail Service and their explanation?

Answer

The main contact point for all kind of general Mail issues is the Computing Helpdesk that can be reached through `mail.support@cern.ch` (or `Helpdesk@cern.ch`).

Some more specific e-mail addresses can be used for specific subjects:

- Listbox questions (mailing lists):
`listbox.support@cern.ch`
- Calendaring questions:
`service-calendar-pilot@listbox.cern.ch`
- Roaming questions:
`service-roaming-pilot@listbox.cern.ch`
- LDAP questions:
`proj-ldap@listbox.cern.ch`

Web Publishing**Question [463] – Calendar Service**

What are the “calendaring” solutions recommended at CERN?

Answer

There is not yet any official/recommended solution for “calendaring”.

A “Netscape CS&T Calendar” project had started (described at <http://cern.ch/proj-ldap/calendar.html>) but unfortunately IT recommend to discontinue the use of the Netscape calendaring pilot as it has no future with the supplier.

Situation on January 2001:

An alternative is being proposed by IT/IS to the IT management and hopes are to have a new pilot ready in March. Considered solutions enable Unix/Windows interoperability, and one possible one is based on *Microsoft Exchange* server, with both *Outlook* and Web access.

Question [469] – Web Radio/TV: what impact?

What consequences does it have for the network if I occasionally listen to a web radio or tune into a TV programme? Do I seriously slow down my colleagues?

Answer

Clearly, if one person listens to the radio there is not too much traffic. If several people, and the people are all on a shared 10 megabit Ethernet then there can be an effect.

There is also a difference between listening to a particular broadcast and tuning into an all-day music station!

Video, of course, is much more traffic.

Text Processing**Question [473] – dvips and file `ma1r8a.pfb`**

When I run `dvips` (following `latex`) I get the error message;

```
dvips: ! Couldn't find header file ma1r8a.pfb
```

Answer

Nothing has changed with the $\text{\TeX}$ installation since June 1999. This error (which comes up more and more) is due to the fact that users want to include EPS files that were generated on *Windows*, where the Arial font sits in the system, and so is thus easily included. When going to Linux or Mac, this font is not present.

In most cases it is enough to generate an empty file for these font files, e.g., the command `touch ma1r8a.pfb` will create such a dummy file in the current directory. If you are lucky, and the font outlines are actually included as a resource in the EPS file, everything will be all right. If not, the picture will have to be regenerated with the font included, since the message from `dvips` tells you that the font in question is “somehow” referenced.

Conclusion: there is no general solution to that problem. If the font is included in the EPS file(s), you can also eliminate the comment specifying that font resource in the EPS file, so that `dvips` no longer complains, or else generate empty font files as indicated above. The fact that “it works” on another system might just mean that the file in question is indeed available on that particular system (e.g. *Windows*).

Scientific Software

(Nicole Crémel, IT/User Support)

Question [404] – Geant3 and Geant4 documentation

Where can I get documentation on Geant (Geant3 and Geant4)?

Answer

Documentation (manual) for Geant3 is available on paper (from the Computing Helpdesk), or from the Web page <http://wwwinfo.cern.ch/asdoc/Welcome.html>.

The Geant4 documentation is only available from the Web (not on paper) at the URL:

<http://cern.ch/geant4/G4UsersDocuments/Overview/html/index.html>

On that page you will find the individual documents in both browsable (HTML, PDF) and printable format (PostScript, PDF).

Question [406] – Compiling/linking with HBOOK on Solaris

I cannot compile programs with HBOOK on WGS (Sun Solaris-2.6 nodes). A simple program compiled with:

```
g77 chf77.f -o chf77 `cernlib packlib mathlib`
```

gives many linker errors:

```
ld: warning: symbol 'pawc_' has differing sizes:
      (file /var/tmp/ccSHE7Xf.o value=0x3a980; file
      /cern/pro/lib/libpacklib.a
      (hbook1.o) value=0x9c40);
      /var/tmp/ccSHE7Xf.o definition taken
etc...
Undefined      first referenced
 symbol        in file
__aintf        /cern/pro/lib/libpacklib.a(dztyp.o)
__s_stop       /cern/pro/lib/libpacklib.a(zend.o)
etc...
```

Answer

There are two issues in your compilation procedure:

1. The warnings about changing the length of the labelled common. You probably should not do it, but the linker usually does the correct thing and takes the longer one.
2. You are linking with g77 and Cernlib (and consequently HBOOK) is built with Solaris f77, so you get unsatisfied externals for the f77 run-time library. You should really compile and link you application with f77. Otherwise, you must specifically include the f77 runtime libraries at the end of your link sequence.

A more general remark is that using the portable CERN `hepf77` command from ASIS (see <http://consult.cern.ch/qa/1582>) will get you the correct compiler (and options), used to build Cernlib, on all our systems at CERN.

Question [408] – Engineering software package licences

I cannot obtain a licence for a PaRC application and I would like to know who has all the available licences.

Answer

Problem licenses with CE application can happen when users unknowingly take licences for a long time - e.g. either because an xterm session may have crashed - or because they have iconised a window. Nevertheless, some licences will also be used for long times by a single user, specially for long calculations on PaRC.

A new tool can help users - specifically in the case you cannot obtain a licence and would like to know who has all the available licences. This tool is available under Unix and Windows.

Under Unix, it is called `cestat`: it will be in your path after you have typed `parc_apps` (alternatively the full path is: `/afs/cern.ch/project/parc/bin/cestat`).

On Windows (95 and NT) the same tool is available under:

```
Start menu -> Computer Aided Engineering ->
Other tools -> Application Statistics
```

It should be quite intuitive to use and there is a small on-line help.

First you should choose a product and click on it. Then a list of features will be displayed. Click on a feature and you will see a list of users currently using that feature. By clicking on a user you will see the information about the person concerned.

Please note: only a subset of applications are included in `cestat` (more will follow). Naturally only those applications for which there is a finite number of licences is included.

Question [458] – FLUKA availability

I would like to use the code FLUKA. Could you tell me how I can get it and on which Platforms it is available?

Answer

You get all information from the Fluka web site, at the URL <http://cern.ch/fluka/>

Question [464] – gcc compiler version

Are there several versions available for the gcc compiler? How do I access a specific version?

Answer

ASIS installs multiple versions of gcc and the `gcc -v` option selects the appropriate version.

Question [498] – PERL modules on LXPLUS

When I try to run PERL with the DBI module on LXPLUS I have an error message:

```
lxplus012 ~[142]> perl -e 'use DBI'
Can't locate DBI.pm in @INC (@INC contains:
/usr/lib/perl5/5.00503/i386-linux
/usr/lib/perl5/5.00503
/usr/lib/perl5/site_perl/5.005/i386-linux
/usr/lib/perl5/site_perl/5.005 .)
at -e line 1.
BEGIN failed--compilation aborted at -e line 1.
```

On DXPLUS, it is fine.

Answer

This is because on LXPLUS your `perl` command comes from the system (`/usr/bin/perl`) and it does not have the DBI module. You should better use `/usr/local/bin/perl` instead.

Unfortunately `/usr/bin/perl` is needed by some system tools and it cannot be removed. `/usr/local/bin/perl` (from ASIS) is the PERL version (common on all architectures) to use in scripts. This is usually not a problem, since the scripts contain explicitly which PERL to use (in the first `#!` line). But when using PERL on the command line, it is necessary (when there is a system version like on Linux) to specify which PERL version to use. Especially if you need modules like DBI, you should take the ASIS version (modules such as DBI which are in ASIS and not in the system).

About the author(s): *Nicole Crémel is working in the User Support group and, with Roger Woolnough, is responsible for the supervision of operations of the Computing Helpdesk. Editing and revising the "Question and Answers" database is a daily task of most members of the group.*

9. User Documentation

Everything related to “computer documentation”: updates to computer documentation distributed by IT (including the User Support reference cards), offerings of the Computing Bookshop, etc.

9.1 User Support Book Catalogue

Roger Woolnough and Jutta Megies, IT/User Support

Situated in Building 513 1-022 (tel. 74050) is the Computing Bookshop provided by the *IT / User Support group*, where CERN users can find computer books and CDs at discount prices. The service is open weekdays from 8.30 to 12.30 or contactable by e-mail to Bookshop@cern.ch.

Books are purchased from some 15 publishing houses and the catalogue offers a selection of documentation aimed at

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

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

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