


CERN COMPUTER NEWSLETTER

CREMEL
Nicole Laurence

RES-S/IT-US



CERN - IT Division
CH-1211 Geneve 23
SWITZERLAND

July – December 2001

CERN-CNLS-2001-003
Vol. XXXVI, issue no 3

Contents

If you need Help	i	5.1 Upgrade of the NICEWWW Service	19
Editorial	iv	5.2 The CERN <i>Search Engine</i>	20
1 Announcements	1	5.3 Network Portable Outlets for DHCP Service Users	21
1.1 Computing Services during the Annual CERN Shutdown	1	6 Scientific Applications and Software Engineering	22
1.2 Recent Changes in Public User Areas	1	6.1 Engineering Applications on Windows 2000, Availability Planning	22
2 Special 35th Anniversary	2	6.2 Monte Carlo Library News	22
2.1 30 Years of Computing at CERN - Part 2	2	6.3 Software Development Tools Service	23
3 Physics Computing	9	7 Desktop Publishing	24
3.1 CERN Council Gives Go Ahead for LHC Computing Grid Project	9	7.1 Installing L ^A T _E X 2001	24
3.2 Highlights from CHEP 2001	10	7.2 XML Applications at CERN	24
3.3 Data Services Group Activities and Plans	12	8 The Learning Zone	25
4 Desktop Computing	14	8.1 A Quick Tour of Compiling, Linking, Load- ing, and Handling Libraries on Unix	25
4.1 Helpdesk and User Support Contacts	14	8.2 Tips & Techniques from Users to Users	35
4.2 Reduced Number of Versions of CCDB Tools (userinfo, userreg, spaceadm)	14	8.3 Questions and Answers from the Computing Helpdesk	35
4.3 Computer Security: Prevent Viruses and Worms	15	9 User Documentation	41
4.4 Password Recommendations at CERN	17	9.1 News from the User Services Bookshop	41
4.5 HEPiX News	18	9.2 User Services Book Catalogue	41
5 Internet Services and Network	19	10 Just For Fun . . .	42
		10.1 Théorème de l'Euro (en français, in French)	42

Merry Christmas!

Happy 2002!

Editorial Information

Editorial Board

Chief Editor Nicole Crémel (CNL.Editor@cern.ch)
Associate Editors Charles Curran
Ludwig Pregernig
Jacques Altaber
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
CERN-CNL-2002-001	Friday 8th of March 2002.

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

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-2001-003/main.ps.gz>. It is possible that, in order to optimize typesetting, the PostScript printable version differs slightly from the master HTML version. In any case the HTML version on the Web must be considered as the reference.

If you need Help

FRONT LINE Services

	Phone	Location	E-mail Address
Computing Help Desk	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: 8:00–12:00 (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 call center 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 “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 SERVICES

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

CENTRAL SERVICES

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

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

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

Service Definition	Name	E-mail Address
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
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 700s/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 (*)

(*) *Activities are now in PS group, and the web site will migrate.*

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

Service Definition	E-mail Address
Computer Security Information	computer.security@cern.ch
Security Alert Reports	cert@cern.ch

Other Newsletters

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/
CASTOR Query Report	http://wwwinfo.cern.ch/htbin/pdp/castor/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	J. Blake, IT	74893	computing.seminar@cern.ch
Technical Presentations	IT Auditorium 31/3-005	S. Cannon, IT	75036	stanley.cannon@cern.ch

Computer Resources Allocation and Divisional Contacts

Computer Resources Allocation Committee (COCOTIME):

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

Secretary: H. Meinhard/IT 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	ETT M. Draper	FI C. Saitta
EST Ch. Delamare	HR M. Buttner	EP A. Bjorkebo	PS A. 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

Dear Nicole Cremel

Re: CERN Computer News Letter (CNL)

Congratulations on being awarded five stars by Emerald Abstracts Computing Coolsites, our indexed reviews of the best of the Web. Only exceptional sites are included in Coolsites and are eligible to display this award.

Coolsites provide the same intelligent access to online resources that we have been providing to more traditionally published material since 1961.

Our specialist reviewers evaluate every site on five key criteria: style, structure, ease of use, quality of information and usefulness to the practitioner. Every site receives a quality rating of between one and five stars and is categorized, with all major aspects of computer science covered.

As an additional accolade, CERN has been added to the Emerald Abstracts Computing Coolsites page.

You can access the review of your site and the criteria against which each site is rated at the URL <http://www.emeraldinsight.com/abstracts/computing/coolsites/index.htm> by following the Computer Applications interest area link...

Congratulations once again on an excellent site. I hope that the addition of this award will raise your profile as effectively as it has other award winners.

Yours sincerely,

Robert Reynolds, Publisher Liaison Executive, Emerald Reviews/Abstracts

The above text is a message that I received recently as CNL editor. Going to the referred links in the "Emerald Abstracts Computing Coolsites" web page you can read:

CERN Computer News Letter (CNL) CERN-CNL-2001-001

We reviewed the whole CERN web site some time ago (and perhaps it is due another look). However, this is a special issue of the CERN computer newsletter to celebrate its thirty-fifth anniversary and take a look back over those 35 years. We have included it in this section because the newsletter always includes a section on physics computing and another on scientific applications and software engineering; in this issue, these are retrospective. In addition, there is a fascinating article by Michel Goossens *A history of scientific text processing at CERN*, which is effectively a fairly comprehensive history of scientific text processing as a whole. We will not list all the other articles here. And, of course, you can always go and look at any of the issues back to 1966. All are available in HTML, most in PostScript and a few as pdf. This is a fascinating series of publications. It perhaps goes without saying that, although the pages are not fancy, they are well designed and clear, as well as being reasonably fast.

It is always pleasant to see that our Computer Newsletter can be appreciated even outside CERN, and also to see that the CNL web edition seems to be well designed ("clear, and reasonably fast").

With this award, we end the "Special 35th Anniversary" on a very positive note. The 35th anniversary CNL was Miguel Marquina's idea. Miguel has left User Support Group to take on other responsibilities within IT division. We would like to thank him for his contributions to the CNL. ¡Muchas gracias y suerte Miguel! from all of us.

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

   
**MERRY
CHRISTMAS!**

... and my best wishes for a happy 2002!

Nicole Crémel, CNL Editor, IT/User Services



1. Announcements

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

1.1 Computing Services during the Annual CERN Shutdown

Tony Cass, IT/TS

Abstract

Annual Christmas CERN shutdown news.

As in previous years, computing services run by IT division will be left running unattended during the annual shutdown. The following points should be noted.

- No interruptions are scheduled for local and wide area networking and the ACB, e-mail and Unix interactive services.
- Unix batch services will be available but without access to manually mounted tapes.
- Dedicated Engineering services, general purpose database services and the Helpdesk will be closed during this period.
- An operator service will be maintained and can be reached at extension 75011 or by Email to `computer.operations@cern.ch`.

Users should be aware that, except where there are special arrangements, any major problems that develop during this period will most likely be resolved only after CERN has reopened. In particular, *we cannot guarantee backups for Home Directory files* for either Unix or Windows users during the shutdown. Any changes that you make to your files during this period may be lost in the event of a disk failure.

Please note that *all the AIS services and databases* run by AS division, such as EDH, BHT, CFU and HR (see also <http://ais.cern.ch>.) *will be unavailable* from 12:00 on Friday the 21st of December 2001 until 08:00 on Monday the 7th of January 2002.

Please remember to *shutdown and power off any equipment in your office* which is not foreseen to be used during the annual closure before you leave for the holiday.

1.2 Recent Changes in Public User Areas

Roger Woolnough, IT/User Services

Abstract

As you may have noticed the two public user areas in 513 and outside the User's Office have noticeably improved.

Public PCs in user areas

At the so-called *Internet Cafe* outside the User's Office four new Windows 2000 PCs with flat screens were installed in July. A further three similar machines were installed at the end of September and the whole area has been restructured; the two X-terminals remain. Although the area is covered by a web camera, the machines are now totally encased in locked cabinets. This also prohibits the use of floppies or CDs.

In building 513 all five PCs available have been loaded with Windows 2000.

Network Sockets for Portable

As you may have noticed in building 513, two new tables have been installed opposite the Macs in the User Area. Adjacent to these are network sockets for portable machines: detailed instructions are available. There are electric plugs

and cables for the network. Users who need plug adaptors or cables should ask the Computing Helpdesk, where spare network cables and electric adaptors are available, against a signature on the forms provided. Outside prime hours the operators have access to this and can help users during nights and week-ends.

Similarly, in the User Area next to the User Office, there are two free network sockets dedicated to portable machines: instructions are available for connections.

Recommendation

We ask users to take care of these machines and to leave them in a state in which they would wish to find them. We also remind users that logon to these machines binds them to the Operational Circular No 5 for CERN computer rules, available at the URL <http://cern.ch/ComputingRules/>.

Any suggestions for improvements may be sent to the author or to `User.Relations@cern.ch`.

About the author(s): *Roger Woolnough is deputy group leader of IT/User Services group, looking after the Computing Helpdesk, the Computing Bookshop and the User Areas.*

2. Special 35th Anniversary

A special chapter for authors who have “missed” the CERN-CNL-2001-001 edition to contribute some historical articles to celebrate the CNL’s 35th birthday this year.

2.1 30 Years of Computing at CERN - Part 2

Paolo Zanella, Former Division Leader of DD

Abstract

This is the second and closing part of a paper written by Paolo Zanella in 1990.

Note that when the word “present” is used, it refers to “1990”.

Transistor machines and data links

The arrival of the 7090 marked the end of the vacuum tube computer. The old Mercury’s Autocode service was stopped and the venerable machine was connected directly to a sonic spark chamber experiment at the CERN PS (Missing Mass Spectrometer). It continued to collect data on-line (one event of some 500 bits per burst) for another couple of years. But the second-generation machines were much faster and much more reliable. The mini’s appeared on the market and CERN bought its first transistorized minicomputer, the SDS 920, in 1964 to be used on-line to another acoustic spark chamber experiment. It had a core store of 4 Kwords (24 bits), paper tape input/output, two magnetic tape units, a printer, a DMA, an interrupt system... and an assembler. It was capable of reading 50 24-bit words of data per event and recording 12 events per burst on magnetic tape while monitoring the experimental equipment. When it detected a faulty device it alerted the skeptical physicists on shift. It took some time before the on-line system could establish its credibility.

The rest of the 60’s saw the decline of film and its replacement by digital events directly acquired by computers. The minis spread all over the experimental floor and by the end of the decade one could count some *50 computers at CERN* (mainly from Hewlett-Packard, DEC and IBM). Data were recorded on tape reels and rushed to the computer center to obtain a quick feedback-using special “*priority cards*.” This “pony express” type of traffic became known as “bicycle-on-line”. The pioneering connection of the Mercury to an experiment had required the design and construction of a *data link about one kilometer in length*. This marked the beginning of the new era of data communications at CERN.

From then onwards, systems for the transmission of digital data have been successfully developed at CERN. From FOCUS (a CDC 3100 based facility for remote access to, and interactive use of the central computers) to OMNET (a network of PDP-11’s clustered around a CII 10070 providing an on-line data acquisition support to the Omega and SFM spectrometers), from CERNET to our current extremely complex system of networks, CERN has been constantly a pioneer in building and using computer communications. From CERNET and its 2.5 Mbits/sec links we learned an impor-

tant lesson, namely that one could work from one’s office or experiment without moving all the time to the computer center. Files could be sent and jobs remotely executed using the CERNET facilities. Recently, the exploding needs of the HEP community have brought high-speed data communications into the limelight. CERN has participated in advanced projects on satellite communications (e.g. STELLA), on LANs and fiber optics.

Today, general purpose networks, LANs and WANs, protocols and standards are part of the experimentalists’ life. We should not forget, however, that the beginnings were slow and difficult.

The epic discussions which accompanied the birth of CERNET showed a deeply divided community reluctantly agreeing to make the initial investments. It wasn’t until the end of the 70’s that the strategic importance of data transmission was fully realized by the most advanced users, the rest joining in during the 80’s.

Centralization and decentralization

The second half of the 60’s were also marked by the introduction of a large central system, the CDC 6600, designed by computer pioneer Seymour Cray. New buzzwords were added to the fast expanding computer jargon, such as *multi-programming, peripheral processors, parallelism*, etc. The word length jumped to 60 bits, the number of registers in the cpu became respectable, instructions were fetched ahead, and huge expensive disks appeared, strong enough to be later recycled as tabletops. Most important, we started talking nanoseconds.

The introduction of such a complex system was by no means trivial and CERN, as many other sites, experienced one of the most painful periods of its computing history. The coupling of unstable hardware (we installed “Serial Number 3”, a pre-production series machine) and shaky software resulted in a long traumatic effort to offer a reliable service. To give an idea of the situation, suffice it to say that CDC had to cancel their SIPROS operating system and CERN had to write large portion of system software to be able to use the machine. In the meantime, an emergency service was set up using more conventional CDC machines, such as the 3400, the 3800, and later the 6400.

Eventually the new hardware and software were fully debugged and the 6600 started a long productive career which ended in 1975, 10 years after its installation. But the growing needs of an experimental program, increasingly based on

electronic detectors, were calling for another big quantum jump in capacity. In 1972, CERN installed a CDC 7600, the most powerful machine on the market, some 5 times faster than the 6600. It was front-ended by two big computers from the same manufacturer, namely a 6400 and a 6500 equipped to manage all the I/O traffic (local or coming from remote input/output stations) and to serve the *first on-line time-sharing users*.

Once more, users and service providers alike, had to go through a very difficult running-in period. The system software was again late and inadequate. In the first months the machine had a bad ground-loop problem causing intermittent faults and eventually requiring all modules to be fitted with sheathed rubber bands! History repeated itself painfully and finally the machinery started crunching an enormous amount of data, and it did it for the longest time of any central computer in CERN until it was turned off in 1984 after over 12 years of service.

In spite of their running-in difficulties, the 6600 and the 7600 were indeed remarkable machines. For 20 years (1965-1984) they played a leading role in computing for High-Energy Physics at CERN as well as at many other laboratories. There was simply nothing comparable on the market. No other machine was so advanced and so fast. The speed of the 7600 processor designed in the 1960's (Oust over 10 Mips) was unbeaten throughout the 70's and the early 80's. Even today's fastest scalar machines in the CERN computer center are only two to three times faster.

The lessons from the early 70's made people reflect on the vulnerability of big, complex systems. One abandoned the idea of connecting film scanners, let alone experimental devices, to a large central computer. The trend towards *decentralization and separation of functions* was irreversibly started. This was supported by the timely arrival of integrated circuits, microprocessors, and powerful minis, which invaded CERN in the late 70's. This vigorous push towards smaller and cheaper computers has continued without interruption during the 80's bringing the current computer population on site to tens of large machines, hundreds of minis, thousands of personal workstations, and an unknown number of *intelligent* micro-devices unaccounted for in any census, each capable of storing, moving and processing data much faster than any first-generation computer. This trend has, however, not reduced the need for central services. Large expensive data handling systems have continued to be operated centrally. Networks, databases and file systems have been managed from the computer center. Last but not least, the computer center has evolved into a computer and communications competence center, where knowledgeable humanware provides a most valuable interface to that soft and hard world of bits and cables.

Multivendor systems integration

In the meantime a significant trend-setting event took place in the computer center, namely the comeback of a large IBM system, the 370/168, in 1976. This computer was not meant to replace any of the CDC machines but to co-exist with

them. At the time it seemed rather risky to manage two large systems from two competing manufacturers, installed side by side in the same room. There were obvious technical and psychological problems which needed imaginative solutions, a different service approach and a new management style. Faced with such challenges CERN found itself leading the way towards the currently prevailing computing environment made of an interconnected set of heterogeneous systems.

The IBM 370/168 brought to the computer center *the silicon chip, the 8-bit byte, the hexadecimal number system*, virtual memory, cache memory, a robotized mass storage system, a set of modern magnetic tape and disk drives, a 19000 lines per minute laser printer, WYLBUR, and above all it demonstrated that complex computer hardware could work reliably!

It will also be remembered, long after its decommissioning, as "the CERN unit" of computing capacity corresponding roughly to 3 Mips (Millions of instructions per second) or to 4 DEC Vups (VAX units of performance). It is amazing to think of this big machine as a major component of the computer center, supporting hundreds of users at the end of the 70's, and to compare it with small, modern personal workstations easily surpassing it in power, in functionality, and providing a user-friendly interface on top.

At the turn of the decade CERN had become a well-known computing outfit where one could see powerful CDC and IBM systems serving a growing number of time-sharing users, but also *talking* to a large and fast increasing population of minicomputers connected via CERNET. The stage was set for playing the role of the *prestigious customer*, and the computer center did not miss the opportunity to become one of those places where all the major computer manufacturers would *rather not take the risk of being left out*...

The harmonious growth of the 80's

In the 1980's large systems and the computer industry in general have undergone a process of Accelerated evolution. Today the profusion of different classes of computers indicates a very rapid *expansion due to the technology and to a certain specialization*. Terms like microcomputers, personal computers, professional workstations, minicomputers, superminis, minisupers, mainframes and supercomputers reflect the fact that problems can now be solved with systems not necessarily large, or so expensive that they have to be shared with other users. The advent of distributed computing services and the continuing advances in technology are making the classification of systems ever more difficult and imprecise.

During this period the computing environment at CERN has undergone a spectacular expansion inside and specially outside the computer center. The 80's have been characterized by frantic cabling of office buildings as well as laboratories and experimental areas. Integration has been the name of the game and standardization has emerged as an obvious corollary. The scene is no longer dominated by one or two high-end machines. Mainframes are part of a network of hundreds of shared computers and thousands of powerful, single-user workstations. If anything dominates this landscape it is the

interoperability, the communications and networking aspect rather than the performance of any particular machine.

Recent evolution of high-performance computing

The high-quality IBM hardware served first with a WYL-BUR sauce and then with a VM/CMS flavor was an undeniable success and the 1980's saw a number of IBM systems succeeding each other ever more rapidly and supporting a population of users which has grown by a factor of ten in the last ten years! The individual mainframes appearing on the CERN stage had IBM numbertnames such as 3032, 3081/D/K, 3090/200/400/600E. They were all based on the same architectural principles, and they offered a smooth growth from the unitary value of the 168 to the power of the current 3090/600E evaluated at 39 units. The main memory size underwent an even more impressive change from the initial 4 Mbytes to the present 256 Mbytes, not to speak of caches and another 256 Mbytes of solid state extensions. But the picture would be incomplete without mentioning the contributions of IBM compatible machines, such as the Siemens 7880 and 7890S which were run under the same operating systems as the IBM mainframes (MVS and VM) and in a way practically transparent to the users. The Siemens 7890S is currently providing some 13 units of capacity and it has set new records of reliability on-site.

In the meantime the CDC line was upgraded in several steps through the replacement of ageing machines by more modern ones, e.g., the Cybers 170/720/730/835/875, until the mid-80's when it was suddenly discontinued. The main reasons were financial and it is true that the computer center had to take its share of the sacrifices required by the construction of LEP. It is also true that this type of machines suffering from an unstable operating environment and faced with tough competition of rival services, was gradually losing popularity in the HEP community.

Exit CDC, enter Cray. Early in 1988 a *Cray vector supercomputer (Cray X-MP 48)* was installed both to increase the center's strength in view of the large amount of LEP data expected, and to see if this type of number cruncher with powerful scalar and vector capabilities could be efficiently exploited by the imaginative HEP community and hence reduce the gap between needed and provided computing capacity. It is clearly too early to see whether vector architectures will play a significant role in the future of computing at CERN. In order to explore as seriously as possible this alternative computational approach, six Vector Feature (VFs) devices have been added to the six processors of the IBM 3090/600.

Since 1982 there is a fourth manufacturer in the computer centre: Digital Equipment Corp. (DEC). Unlike the other vendors, they found their way to the top after having spread hundreds of PDP's and VAXes all over the site. Physicists were clearly fascinated by the user friendliness of the hardware and of the software. Digital made neither *black boxes* nor... *blue boxes*. They could be opened, connected to experimental equipment, and coexist with home-made elec-

tronics in the counting room. The fierce competition (e.g., Hewlett-Packard, Norsk Data, CII, IBM, etc.) had difficulties to survive the aggression of the 32-bits VAX family propelled by VMS and DECnet. After almost all the experiments had been contaminated (or "VAXinated"), and even the LEP machine's database and CAD system were installed on VAXes; a DECcluster was set up in the central computer room to offer a number of services initially reserved to the LEP experiments. The current central VAX system consists of an 8800, an 8700, three 8650 and an 8530. The total VAX count at CERN is well over 300 and still growing.

Two more success stories have marked the last few of the past 30 years, namely the *personal computer and the workstation*. The first are gradually invading the buildings and are used for a variety of administrative and office tasks, e.g., text processing, document handling, electronic mail, preparation of presentation material, spreadsheet and database work, etc. Two types of personal computer are present at CERN and account for well over 2000 systems: the ubiquitous *Apple Macintosh* and the *EBM or EBM-compatible PCs* (e.g., Olivetti), with the first type currently growing much faster than the second. For the record, the first IBM PCs appeared on-site in 1983.

As to the workstations, they are, and will be even more so in the future, a major factor of the computing environment. Apollo workstations connected into local area distributed systems are already over 100 and growing fast. There are also several clusters of VAX workstations which have proved particularly useful for program development. The power of these machines is already quite impressive and is increasing quite dramatically every year, while their price is dropping due to technology and competition. It is easy to predict a very healthy future to the advanced workstation as part of a distributed system providing access to file servers, communications servers, print servers and compute servers. While one is learning how to solve the technical and managerial problems posed by these new systems, the number one challenge is how to face the extremely rapid obsolescence of such devices and how to plan and finance their continuous replacement.

To conclude this quick walk through the complex evolution of computing at CERN one should look at Fig. 1, which shows the growth of the central computers' usage from the Mercury to the current situation, expressed in CERN cpu units. It is true that cpu units do not tell the full story. If we, however, plot the number of tapes in the vault (a few hundred thousand to date), or the disk capacity against time, we get even steeper climbing curves. It is worth noting that this exponential progression has been necessary to support a rapidly increasing population of users. Thirty years after the first attempts at running an Autocode program, there are today over 6000 registered users of the CERN computer centre depending more and more on its facilities and services. It is also interesting to note that for the last 20 years this growth has been achieved at almost constant budgets and staff numbers.

From general purpose to HEP-specific computing

During the 30 years under consideration a few major trends have been driving the evolution of computing at CERN and, as a matter of fact, everywhere else. Among these forces there is a strong attraction towards decentralization and separation of functions. This is probably one of the causes and not a consequence of the extremely diverse market offer.

There is a definite tendency to specialize systems, associating them to a particular service, e.g., CAD, LEP Database, Microprocessor support, and integrate them in the global informatics environment. This gives better service, better reliability and easier management. The on-line computing function is no exception and it is spreading over a number of machines each assigned to a particular set of jobs. A similar tendency can be observed in the accelerator control systems.

At CERN we have had a tradition in the design of special hardware and software. We seem to believe that our problems are very special and that we can solve them better than anybody else. It is fair to say that in some cases we may even be right!

As far as the hardware is concerned, we started by building a whole family of film measuring and digitizing equipment and when the film was replaced by electronic detectors we built all sorts of *special processors, programmable fast trigger devices (ESOP, XOP) and emulators (Mice, 168/E, 3081/E)*. The emulators have been a real success in the last decade, both on-line as a high-level trigger technique and off-line as cheap event processors. We also designed and built the first communication systems (data links, FOCUS, OMNET, CERNET, etc.) well before one could buy network components off the shelf. The fact that one can now buy high-functionality building blocks merely shifts the integration job to a higher level.

One can tell a similar story for the software side. After the initial original contributions to basic system software (paper tape drivers, assemblers, compilers, operating systems, program libraries, data base management systems, etc.) the effort has shifted to higher levels of system integration and to specific HEP problems, e.g., HEPVIM, ZEBRA, PAW, GEANT, etc. Moving closer to experiments, the evolution of computer applications has involved an ever larger fraction of the participating physicists. The computer wizard of the 60's has been replaced by the local guru of the 80's, but out in the field computing can still trigger religious wars!

The landscape at the start of LEP

At the beginning of the LEP era the CERN Computer Centre offers a wide range of computing and communications facilities. The main systems for physics based on IBM, Siemens/Fujitsu, Cray and Digital mainframes account for almost 90 CERN units of installed scalar capacity (plus a substantial vector capability), roughly 100 000 times the power of the Mercury. Taking into account all the machines existing on-site including the workstations and the personal

computers, the overall computing capacity has expanded by almost six orders of magnitude in 30 years.

Data communications at CERN cover electronic mail, file transfer, remote job entry and terminal access services. The ageing CERNET is still in use while an expanding set of Ethernet segments bridged together cover the whole site. 2000 terminals are connected to a Gandalf terminal switch called INDEX which has been operated for the last fifteen years and still allows users to work on any one of a large number of systems. The number of personal computers and workstations connected to Ethernet, the LEP token ring or any other local area network, is already well over 2000 and growing rapidly. A number of protocols are used over the local networks, e.g., CERNET, TCPAP, OSI, Digital DNA and IBM SNA.

CERN is linked to universities and research laboratories all over the world via leased lines and the public X.25 network. A variety of services are provided over these lines. CERN acts as the central site for X.25 and DECnet networks used by the HEP community and as the Swiss node of EARN and EUNET. Currently the prevailing line speed is 64 kbps, but an urgent upgrade to 2 Mbps circuits has been advocated and is just starting.

At the experiments, data are filtered in real time and collected on tape by complex arrays of processing devices interconnected by high-speed bus systems, e.g., FASTBUS, VMEbus. Several techniques are applied in successive stages, e.g., fast special processors, dedicated microprocessors, emulators, conventional mini's, workstations and midrange computers. A very important milestone, UA-I (1983) can be considered the first of a new generation of experiments, making extensive use of computers and microprocessors (over 300) on-line to the detectors.

The selected, compressed and recorded experimental data are then submitted to further processing off-line at CERN as well as at the physicists' home institutes. Large scale special programs are developed for this analysis job. Software engineering techniques are beginning to be adopted by the larger collaborations. It is clear that the design and production of software can no longer be regarded as the creative *outburst of some gifted individuals*. Analysis chains are made of millions of lines of codes to be executed on millions of events stored on tens of thousands of tape cassettes. The off-line activity has become a major management job. It also includes the simulation task, another very important and fast growing aspect of computation.

The final analysis stage makes increasingly use of the powerful graphics capabilities of professional workstations. Complex data patterns are visualized as 2, 3 or multi-dimensional images. Recently, the "Physics Analysis Workstation" Project (PAW) has produced a distributed application which takes advantage of the interactive graphics power of modern workstations e.g., Apollo, DEC, and the number crunching and large storage capabilities of central mainframes. The introduction of interactive graphics techniques at CERN dates back to the 60's when we were using large and expensive CDC first generation displays to visualize physics events

and mathematical functions (e.g., GAMMA, SIGMA). The first large scale project to exploit interactive graphics was ERASME, which used DEC computers (initially PDP 10 and 11's, and later, VAXes) to scan and measure bubble chamber films. *The expertise built up throughout the 60's and 70's* was then used in the 80's to design the modern event visualization and analysis systems (e.g., MERLIN developed for the UA-I and UA-2 experiments), which are now widely used in experimental high-energy physics.

Three decades of expanding computer applications have left behind a huge amount of software which after proper selection, sorting, and documentation has formed the basis of an invaluable treasury: the *CERN Program Library*. Continually updated and distributed worldwide, the Library is at the same time an active tool, an extremely useful collection of numerical procedures, and a unique repository where all the experience accumulated over 30 years of computing at CERN is conserved.

After 30 years of continuous development the field of computing is still changing very fast. The attention has been shifting from input/output to cpu speed and main memory size and then back to input/output and peripherals, from algorithms to programming languages and then to operating systems, from data links to networks and to distributed systems. The current problem is all of the above plus *large scale data handling*: we need dramatic improvements in data storage and transmission systems. The scale is new but the problem was already there with the Mercury...

New challenges and exciting opportunities appear every year. The problems of high-energy physics have never been so complex, but never before technology has been so rich and generous in proposing solutions. Life is hectic for planners and service providers who can hardly match user's expectations and have a hard time keeping the environment stable while replacing hardware and software at an unprecedented rate.

The landscape is totally changed. The computer is no longer a technical curiosity. Its impact has reached levels far beyond imagination. Hardly anybody is unaffected. Yet the field is moving too fast to be really close to maturity. Success or failure still depend strongly on people. While expertise is spreading and most physicists, engineers and administrators are more comfortable with computers, our systems are vulnerable and depend on a few key experts. This is after all reassuring. Contrary to popular scenarios of the 50's, and in spite of their shortcomings, human beings are still in the driver's seat.

The next 30 years

In 30 years time the *CERN Computing School* will celebrate its 50th anniversary. If I had to pick one thing likely to still be alive 30 years from now I would choose *FORTRAN*. It is as safe a bet as to predict that everything else is going to change.

Of course there will be limits to growth. The number of users will level off one day since we are approaching a predictable

ceiling. The number of personal computers and workstations, however, has still plenty of room for growth given that we are far from the limit of 2-3 systems per employee, already experienced in "high-tech" industries. For sure, the financial limitation is going to be felt for quite some time.

Technology promises to continue to advance at the current rate for the foreseeable future. By the time when one will no longer be able to squeeze any more juice out of semiconductors and magnetic recording, optical techniques will come to the rescue. So, it should be possible to continue for a while to get more *hfips*, *Megapixels*, *Giganops* and *Terabytes*. And surely, we will know what to do with them, if anything at all will be left for the user. The long awaited standardization of communication protocols, operating systems and application environments, added to the well-known appetite for resources of graphic layers, friendly user interfaces, encryption, file and database systems, will probably keep the user in a state of constant *unhappiness*...

The question is whether there will be any dramatic changes. I mean something even more dramatic than moving to a 16-bit byte, seeing a complete set of ISO protocols replacing TCP/IP, or watching physicists using *C* and *UNIX*... Will people stop programming, either because every possible program has been written or because neural networks will replace programmable circuits? Will the myth of the almighty central brain be revived by some super massive parallel engine leaving only text processing, Mac paint and video games to the personal play stations? When will a super friendly human interface entertain the user in his own language and provide him with all the information he wants without having to type and mouse? When will we all be working from home? Will the computer ultimately collapse to an intelligent pinhead, whilst data will expand to fill any available space?

Waiting for all this to happen, I invite the young generation of computing physicists to a short pictorial tour of the past: the selection of "historical" images presented at the end of this paper should help understand and remember those very intense and exciting, first 30 years. I realize that this short account does not cover many important aspects of the history of computing at CERN (e.g., Accelerator Controls, Computing in Theoretical Physics, Office and Management Information Systems, etc.), and that it only mentions very briefly many developments, which would deserve a more extensive presentation. The rich specialized literature existing in the laboratory should help satisfying the legitimate curiosity of the history-minded student.

Technical appendix

CDC 6400/6500 [1967-1980]

The 6400 was architecturally similar to and compatible with the 6600, but less powerful (40% of a 6600). The twin processor version of the 6400 was called 6500 and CERN upgraded its 6400 to a 6500 in 1969 to be used as a backup machine, and later as one of the two front-end to the 7600 (the other being another 6400). They were eventually replaced by

a pair of Cyber 170 machines (720 and 730) at the end of the 70's.

CDC 7600 [1972-1984]

Designed by Seymour Cray, the CDC 7600 was an astonishingly compact and elegant machine. It came with 64 Kwords (60 bits) of small core memory with a 275 nsec read/write cycle time (SCM) and 512 Kwords of large core memory with a 1760 nsec cycle time (LCM). Both were still ferrite core memories. Cpu clock cycle = 27.5 nsec. Arithmetic power: 1 10 nsec add and 13 7.5 nsec multiply. An Instruction Word Stack of 12 60-bit registers contained prefetched instructions for faster execution by avoiding frequent memory references. Capacity of the 7600 estimated at 3.3 CERN units. The 7600 had 15 Peripheral Processors. It was run under the 7600 Scope 2 operating system. Front-ended initially by a 6400 and a 6500 and at the end by a Cyber 170/720 + 730 pair running NOS/BE, and sharing a 512 kwords ECS (Extended Core Storage).

Like the 6600, it suffered from teething problems (hardware and software) during the first two to three years. The mean time between failures requiring a deadstart of the combined 7600 + 6400 system was 3 hours in 1973, 7 hours in 1974 and still only 9 hours in 1975 three years after delivery. A more acceptable 16.3 hours MTBF level was reached in 1977 and the machine continued to improve its reliability at the price, however, of considerable effort. The 7600 has been operated for 12.5 years. During its exceptionally long life it has processed a total of 6 681 378 jobs and delivered to its users 61321 cpu hours (over 60070 of total time)!

IBM 370/168-3 [1976-1982]

The 168 was delivered with 4 MBytes Oater expanded to 5) of semiconductor memory, a 16 kBytes cache, 4 channels (later 7) and a high-speed multiply unit. The cpu cycle was 80 nsec. It was a virtual memory machine. Although it became "the CERN unit" of physics data processing power, its strengths relative to the CDC 7600 varied strongly, being e.g., weaker in floating point calculations and stronger when *large memory* was needed. The hardware reliability, including tape (9 tracks, 1600 bpi) and disk units, was a significant improvement on previous systems for HEP computing.

It ran the MVS (Multiple Virtual Storage) operating system with the JES2 Job Entry subsystem for batch work, and the WYLBUR/MILTEN terminal system for an eventual maximum of 200 concurrent users editing on line-mode terminals. The exceptionally well-designed and friendly user interface of WYLBUR made it a firm favorite with the users, despite its inherent limitations of having neither full-screen working, nor allowing fully interactive computing.

IBM 3032 [1978-1981]

The 3032 was a re-packaged 370/168, lacking only the high-speed multiply unit. It was delivered with 6 MBytes of main memory, and 6 channels which were, in contrast to the 168,

incorporated in the mainframe. Its performance was slightly weaker (10010) than the 168 for typical CERN production work, but its higher memory and I/O capacity allowed it to be put to good use as a "front-end" machine. The system software was the same combination of MVS/JES2/WYLBUR run on the 168.

IBM 3081 [1981-1985]

The 3081-D replaced the 3032 in September 1981. It came with 16 MBytes of main memory, 16 channels and 32 KBytes of cache memory per processor. This was the first IBM "dyadic" machine, having two "tightly coupled" processors, which, unlike the previous MP systems, could not be separated into two independent systems. It was also the first IBM machine to use the "Thermal Conduction Module" (TCM) packaging, also used later in the 3090 family of machines. The cpu was built using TTL chips.

The installation of the 3081 came as a great relief to the previous 168 + 3032 system which was saturated to the point of serious performance degradation. The CERN benchmarks rated each of the processors of this D-model as 1.9 CERN units. The machine was upgraded to a model K in 1982, giving each cpu a power of 2.4 CERN units. The 3081 was subsequently upgraded to 24 MBytes of main memory and 24 channels. The performance was further increased by 10070 in 1984. Finally, it was sold in December 1985, and was replaced by an IBM 3090-200.

Siemens 7880 [1982-1985]

Made by Fujitsu as the M200, using technology developed together with Amdahl, and sold in Europe by Siemens, this IBM compatible machine had 12 MBytes of memory and 16 channels. The single cpu had a power of 2.4 CERN units. The 7880 was acquired to replace the IBM 168. It ran the MVS/JES2/WYLBUR system together with the IBM 3081 in a manner transparent to the users.

CDC Cyber 170/875 + 835 [1983-1986]

Acquired to replace the 7600 system these computers had respectively two and one cpu while each had 1 Megaword of memory (75 nsec memory cycle). The 875 had a cpu clock cycle time of 25 usec. They shared disks and tape units and ran under the CDC NOS/BE operating system. Compared to the 7600 + 720 + 730 complex running under Scope 2.0, this system was rated at least twice as powerful, each 875 cpu providing 3.5 CERN units and the 835 adding another 0.6 units. Compared to the IBM and Siemens/Fujitsu machines, the reliability of this system was somewhat disappointing and in 1984 CDC had to take some corrective measures in order to improve the service. Both machines were decommissioned in October 1986 and their departure marked the end of a 22 year long collaboration with Control Data.

Siemens 7890S [1985-present]

Made by Fujitsu (as the M382) and commercialized in Europe by Siemens, this twin-cpu IBM compatible machine is an early version of a speeded-up series and as such has a capacity of around 6.5 CERN units per cpu which is very similar to the power of the 3090-200. It has 64 MBytes of main memory in 256 Kbit *dynamic RAM* chips, and 32 channels, but no expanded storage. The 7890 has a double level cache for each cpu, the first level consisting of a 64 KBytes of 5.5 nsec, 4 Kbit/chip static RAM, and the second level (global buffer) consisting of 512 KBytes of 16 nsec, 16 Kbit/chip static RAM.

The machine is capable of running IBM extended architecture system MVS/YA but not VM/XA because it lacks an instruction introduced later by IBM. Nonetheless it is an impressive machine given its appearance three years before the 3090/200 of comparable speed. At CERN, where it has set new standards of reliability, it shares peripherals with the IBM 3090 system and it has been run under the same IBM operating system (initially MVS/JES2/WYLBUR, and later VM/CMS) until the end of 1988 when the 3090 started operating under VM/XA.

MM 3090 [1986-present]

This is the current flagship of the IBM mainframes. The 3090-200 was installed in December 1985, with 64 MBytes of main memory in 64 Kbit chips, and 64 MBytes of expanded storage. It had two processors, each rated at 6 CERN units, and 32 (later 40) channels. The cache memory consisted of 64 KBytes per processor, and the cpu cycle was 18.5 nsec. The cpu was built in ECL chips. Above all the machine could be considered as IBM's first real scientific processor since the 370/195, as it is optimized for 64-bit word arithmetic.

The machine was upgraded to a model 3090-400E in May 1988, with 256 MBytes of main memory in 1 Mbit chips, 256 MBytes of expanded storage, and a total of 80 channels. In September 1988, four vector facilities (VF) and 16 channels were added, and at the end of 1988, two more cpu's and two VFs were installed, making a full 3090-600E with 6 VFs. Each processor is rated at 6.5 CERN units of scalar power, has a cycle time of 17.5 nsec, and 64 KBytes of cache memory. This model 600 can also be considered as a twin "triadic machine".

The disk capacity has reached 200 GBytes towards the end of

1988 and is supposed to grow very fast in the coming years. Old tape units are being replaced by cassette drives (IBM 3480: 18 tracks, 19000 bpi, holding over 200 MBytes). The operating system has been gradually changed from MVS to VM. Currently the machine runs under VM/XA. The IBM VM service is the workhorse of the CompaWr Centre. This installation is part of the IBM European Academic Supercomputer Initiative, a small number of IBM 3090's equipped with VFs and interconnected via EASInet (currently at 64 Kbps).

DEC VAX 8600/8650/8700/88001 (1985-present)

The central VAX cluster grew out of a set of VAX 11/780s installed in the Computer Centre from 1982 and dedicated to database (ORACLE) and CAD (EUCLID) services for the LEP machine designers.

In 1985, following a major upgrade including the installation of three VAX 8600s, an interactive VMS service was started, limited to the LEP experiments. In 1987 the VAX cluster was further expanded to include a twin-processor VAX 8800 and three 8650s. The VMS service offered to up to 200 simultaneous users was a big success and the central VAX cluster has become a focal point for the VAX/VMS community at CERN. With the addition in 1988 of an 8700 and a 8530 this cluster of DEC hardware has become a non-negligible entity in the Computer Centre with over 3000 registered users.

The largest machine in the cluster, the 8800, has two cpu's each providing about 1.5 CERN units of capacity and a main memory of 96 MBytes. The machines connected to the cluster share access to 18 tape units and over 50 GBytes of disk storage.

CRAY X-MP/48 [1988-present]

The Cray X-MP/48 supercomputer has been installed in January 1988. It has a clock cycle time of 9.5 nsec and it features 4 cpu's, 8 Mwords (64 bits) of central memory, and it is running under UNICOS, a UNIX based operating system. It has powerful vector capabilities, 64-bit registers, 128 Mwords of solid-state backing store, 6 IBM 3480 tape cartridge units and some 48 GBytes of disk storage. Users' access is via IBM VM/CMS and VAX VMS. The total scalar power of the machine is about 32 CERN units. It is hoped that selected applications can be sufficiently vectorized to take full advantage of the vector capabilities of the Cray and thus gain appreciable speedups.



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 CERN Council Gives Go Ahead for LHC Computing Grid Project

Press and Publications Service, CERN

Abstract

Development of the LHC Computing Grid will be done in collaboration with institutes in Member and non-Member States working under a coordinated project structure, whose Project Leader will be Les Robertson. The project's main objectives are to apply Grid and other technologies towards building a set of increasingly complex and powerful world-wide distributed prototypes, in order to ensure a proper foundation for Phase 2 of the project, where the enormous initial production system capable of receiving and analyzing data from the LHC experiments will be built up. The project aims to use the results of projects working on basic Grid technology development, such as EU DataGrid, GriPhyn, PPDG, etc., integrating their deliverable schedules with the prototype deliveries, which in turn will be integrated with the Data Challenge needs of the LHC experiments.

Manuel Delfino, IT Division Leader.

A thousand times more computing power by 2006

The first phase of the LHC Computing Grid project was approved at an extraordinary meeting of the CERN Council on 20 September 2001. CERN is preparing for the unprecedented avalanche of data that will be produced by the multi-national experiments at its new particle accelerator, the Large Hadron Collider (LHC), the most powerful device of its kind ever built.

CERN's need for a dramatic advance in computing capacity is urgent. As from 2006, the four giant detectors observing trillions of elementary particle collisions at the LHC will accumulate over ten million Gigabytes of data, equivalent to the contents of about 20 million CD-ROMs, each year of its operation. A thousand times more computing power will be needed than is available to CERN today.

Nearly ten thousand scientists at hundreds of universities around the world will group in virtual communities to comb the data, searching for new physics. The strategy CERN, together with its international partners, has adopted to analyse and store this unprecedented amount of data is the coordinated deployment of Grid technologies at hundreds of institutes which will be able to search out and analyse information from an interconnected worldwide grid of tens of thousands of computers and storage devices.

The LHC Computing Grid project has two phases. The first phase activated in 2002 and continuing in 2003 and 2004 will develop the prototype equipment and techniques necessary for the data-intensive scientific computing of the LHC era. In 2005, 2006 and 2007 phase 2 of the project, building on the experience gained in the first phase, will construct the production version of the LHC Computing Grid.

Phase 1 of the project will require a 30 million Swiss Francs investment at CERN, which will be financed by individual contributions by CERN's Member States and major involvement of industrial sponsors. Over 50 positions for young IT professionals will be created in Phase 1 of the LHC Computing Grid project offering unique training possibilities. Young computer specialists in CERN's Member States, will be at the heart of the development of new data-intensive science computing and develop techniques and knowledge which will be invaluable.

Significant investments are being made by participants in the LHC programme around the world, particularly in Europe, the US and Japan. CERN and its partners will coordinate their activities to form the first "virtual organisation" to adapt Grid technologies to a giant data-intensive, world-wide, computing problem, but the challenge of analysing and storing very large quantities of data will be faced by governments, commercial companies, corporations and institutes in the near future.

The green light for Phase 1 of the LHC Computing Grid project starts another exciting phase of development in CERN's Large Hadron Collider project. The new accelerator will provide collisions at higher energies than ever before. The frequency of these collisions is far beyond anything which physicists have experienced and both the accelerator and the four giant experiments depend on the most advanced technologies available to scientists today. Computing will be the key link between this high-tech hardware and particle physicists, enabling them to make new breakthroughs, the immediate goal being the discovery of the Higgs Boson. Searching for a needle in a haystack is a legendarily difficult job, faced with the tremendously high collision rate in LHC, uniquely identifying a collision where a Higgs was produced will be equivalent to finding a needle in 20 million haystacks!

CERN is confident that the LHC Computing Grid project will give the international community of physicists working on the LHC the advanced infrastructure necessary to solve this seemingly impossible puzzle. At the same time it will provide a catalyst for Grid development around the globe.

3.2 Highlights from CHEP 2001

Vladimir Bahyl, Harry Renshall, Gabriele Cosmo, and David Myers

Abstract

A short summary of some the presentations given at this year's Computing in High Energy Physics (CHEP 2001) conference in Beijing, as reported by some CERN/IT participants.

Systems and Software (Vladimir Bahyl, IT/PDP)

Commodity Hardware and Software

In the Commodity Hardware and Software track, there were several reports on the experiences with running big farms from CERN, DESY, FNAL, BNL, IHEP and KEK. Here are the aspects that were common for most of the talks:

- it seems that Linux with dual-CPU PCs is the current trend all over the globe
- while DESY's configuration is based on SuSE, most of the other are using RedHat
- as for the spread of rack mounted solutions, they are used in the US on considerable scale in contrast to CERN and DESY where there is no problem with space
- concerning the question of the size, CERN's farm is one of the biggest running, even though not always using the smartest solutions

The only common technology used seems to be AFS: for monitoring, job control and data delivery everybody uses something different (often made at the particular institute).

Other interesting talks

Robert Cowles from SLAC gave a very interesting technology review for Security of the data grids. He discussed possible authorisation and authentication solutions based on PKI. Marcin Nowak from CERN presented a talk on object features of Oracle 9i like object modelling (according to SQL 1999 standard), Oracle Type Translator for C++ and C++ binding via OCCI that could be useful for the HEP community.

Andrew Hanushevsky from SLAC presented a talk where he discussed usefulness of compressed database support in Objectivity for BaBar experiment database. While compression trades disk space for CPU resources, the trade-off may be worthwhile in certain circumstances. However - the compression should only be considered as an option while the proper design of the structures in the database is the key.

Networking (Harry Renshall, IT/DS)

There was a half day parallel session on networking. The driving topics were high throughput wide area performance with five papers, and security with three papers.

Several of the performance talks reported on how to get the best out of existing systems for wide area file transfer over TCP and emphasised that current operating systems have bad defaults notably small TCP window sizes. A combination of optimising TCP window size at source and sink (typically to be more than 1MB), using large file sizes (several hundred MB and above) and transferring in multiple streams gave the best results. A user should today be able to transfer 100 GBytes/day between major HEP sites.

Two talks reported easy to use interfaces which do a lot of the optimisation (the window size is a system parameter) namely "bbcp" (BaBar wide area copy) from SLAC and "gsift" from Globus. There was also a very interesting wide area networking simulator using an intermediate PC as a router reported from KEK.

The security talks did not report many new developments but all sites recognise the increasing importance of access security and will be using appropriate tools from the basic ones such as SATAN and COPS to implementing strong authentication (at FNAL) and virtual private networks (KEK).

The IPV6 tutorial was a reminder of industry plans to expand IP addressing from its current restriction of 32 bits to 128 bits as agreed in 1994. Also included is stateless auto-configuration which will allow to rapidly renumber a site. The first users are expected to be the third generation mobile phone developers in Europe and Asia where there is a shortage of address space.

Grid Computing (Harry Renshall, IT/DS)

There were three busy half day sessions plus some plenary talks. There was a session of experiment or project status reports I did not attend then a session on basic tools and one on progress in middleware.

Steady progress on implementation of the Globus architecture was reported and results of a practical study on dynamic replication strategies within Globus. There was a very detailed practical evaluation of Globus by INFN sites working within the European data grid illustrating their good communications with the Globus team and also a much appreciated talk on a recent analysis of the CMS requirements for the data grid.

Middleware progress talks covered a wide area: integrating grid tools to build a computing resource broker (INFN); sparse query processing models (U. of West of England); parallel ROOT interactive analysis system (CERN); a project for peta-byte scale data intensive analysis (GFARM from KEK); and a grid aware extension of the Liverpool University Monte-Carlo array processor allowing remote job preparation and submission using the Globus toolkit for authentication and communication.

There were several relevant conclusions drawn by the rapporteur of the Grid sessions: activities are still mainly concen-

trated on strategies, architectures and tests. There is general adoption of the Globus layered architecture and basic services. New middleware tools are starting to appear and be used but there are some parallel developments so strong co-ordination will be needed and we should plan carefully the next iteration of Grid middleware development. There is in general good collaboration between the existing EU and US grid projects that must continue. Experiments are getting on top of grid activities. And finally the grid infancy phase begun at the Padova CHEP has now ended at CHEP Beijing as evidenced by the recent creation of the LHC Computing Grid project.

Software Methodologies and Tools (Gabriele Cosmo, IT/API)

In the track for "Software Methodologies and Tools", the need to apply Software Process techniques according to well established and standard procedures was expressed. Software Process Improvement must be part of this and be life-cycle driven.

Concerning OO programming, contributors stressed the importance of utilising loosely coupled components and maximise re-use for the development of applications. In this context, the concepts of "collaborating frameworks" and "abstract interfaces" [1] play a central role. The resulting integration of software components, which comply with this rule, is very easy to achieve. At the same time, optimal flexibility and maintainability of the software is assured.

The importance of considering Software Quality and applying it in the normal development process was touched upon. An interesting tool (Ignominy) for quantifying modularity through software metrics was presented. The tool is currently adopted in the context of the CMS software [2].

For Software Management, the features of CMT (Configuration Management Tool) [3] and SCRAM (Software Configuration, Release And Management) [4] packages were presented and contrasted. Site-specific configuration issues for cross-laboratories distribution of the software remain an area of active work.

The Geant4 Toolkit [5] is becoming the emerging standard in HEP for "Simulation". After almost 3 years from its first public release, the software is now used in production on several experiments (ATLAS testbeam, BaBar, HARP) and other intend to follow in the next year. Comparison studies with experimental data were presented with positive results. A variety of applications outside HEP are making use of it, ranging from medical applications to astro-physics studies, thanks to the low-energy extensions of electromagnetic interactions.

In the track for "Analysis Tools", many improvements and new features were implemented in each of the tools presented at the conference: Anaphe/Lizard [6], IGUANA Interactive Analysis [2], Java Analysis Studio (JAS) [7] and ROOT [8]. The Abstract Interfaces for Data Analysis (AIDA) [1] are putting in a strong bid to become the HEP standard for defining the interfaces in the data analysis domain.

For HEP graphics toolkits, QT for graphical user interface, and OpenGL/OpenInventor for low/high-level graphics are widely used. Python as scripting language is showing increasing popularity among developers.

- [1] <http://aida.freehep.org>
- [2] <http://cern.ch/iguana>
- [3] <http://www.lal.in2p3.fr/SI/CMT/CMT.htm>
- [4] <http://cmsdoc.cern.ch/cgi-cmc/scrampage>
- [5] <http://cern.ch/geant4>
- [6] <http://cern.ch/anaphe>
- [7] <http://www-sldnt.slac.stanford.edu/jas>
- [8] <http://root.cern.ch>

Controls Track Summary (David Myers, IT/CO)

Controls seems to attract much more interest from the accelerator community than from experimental physicists, and the accelerator controls people tend to go to the ICALEPCS [1] conference series, rather than to CHEP. Thus, at CHEP, controls seems to be rather a minority interest which is evident from the modest number of contributions. I suspect that this may be due to compartmentalisation into "Slow Control", "DAQ Control", "Run Control", and so forth, rather than realising that there is a great deal of overlap between these areas which I believe to a large extent could use many similar tools and techniques.

There were eleven papers and posters contributed at this conference, plus a plenary talk on the LHC Joint Controls Project, JCOP, and a live demonstration. Of the contributed papers, several discussed interlock and safety systems, one supervised by LabView; one paper dealt with supervision of calibration; one the supervision of a computer farm; and two discussed applications developed with EPICS. Another contribution mentioned CDEV, which was originally a technology developed at Jefferson lab in order to unify different control systems with EPICS. Five papers mention systems developed using commercial SCADA tools, and a paper from ATLAS discussed how to link purpose-built Run Control with a SCADA system. Finally, Clara Gaspar from LHCb bravely gave a live demonstration of "Partitioning, Automation and Error Recovery in an LHC Experiment" using the SCADA system chosen by the LHC Joint Controls Project.

Thus, the topics covered a wide area, from classical hardware monitoring and supervision, to interlock systems and supervision of computing farms. Like all areas of HEP, controls software has historically been home-made. However, an unmistakable trend is now towards the use of Supervisory Control And Data Acquisition (SCADA) systems. One of the first of these used in HEP, EPICS, was developed originally at Los Alamos. However, just as we no longer write our own operating systems and compilers (which was still the case when I arrived at CERN), the use of home-made tools for building controls systems is slowly being overtaken by the use of industrial software. The crucial developments which permit this are the features of the latest generations of SCADA software which provide scalability to very large distributed systems and support for object-oriented, or device-oriented, development.

This movement towards industrial software tools is, perhaps, the major conclusion for the future: rather than putting our effort into building purpose-designed tools, we should use tools we can buy and put our effort into solving our controls problems. The cost of implementing and supporting a purpose-built tool is no longer commensurate with its added value. Although not a subject from the conference, it is also worth mentioning that in some areas, such as cryogenics, power distribution and safety, complete systems are now being out-sourced to industry.

Finally, although I must admit to a possible bias, I want to mention the progress reported by Wayne Salter on the LHC Experiment's Joint Controls Project. Not only is this showing that commercial components can be used in many HEP applications, but also that it is possible for experiments to reduce their development and maintenance costs by working together.

- [1] International Conference on Accelerator and Large Experimental Physics Control Systems.



3.3 Data Services Group Activities and Plans

Harry Renshall, IT/DS

Abstract

A review of the current activities and plans of the DS group in the services of AFS, backup, managed storage, tapes, tape drives and tape robots. The group was recently formed out of the Data Management section of PDP group.

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

AFS Services

Over the summer months there was a clear degradation in the quality and stability of the AFS services with many clients, mostly running Linux, experiencing access problems whereby directories or files apparently disappeared or clients froze completely. There was certainly a large increase in load at this time as many new clients appeared but there was no obvious single cause as both we and the communications and Linux systems staff investigated. We scheduled an upgrade of the fileserver software as there were fixes for some of the locking situations and this was performed on 14 August involving a total stop of AFS services for 1 hour. There was a noticeable improvement in some of the cases but the real change came on 17 August when a network switch in the computer centre was changed. Since then these services have returned to an acceptable level. The next changes here will be to introduce larger scale scratch space servers as part of replacing old non-Linux systems and deploying next year's capacity. In the longer term we will be looking at using OpenAFS file servers instead of using IBM AFS (we already have many Linux clients using OpenAFS).

Backup Services

IT does not backup the individual desktop expecting local disk to be for scratch purposes only and permanent files to be either on the Windows servers, in AFS or under managed storage, either used from there or copied locally. We backup critical computer centre machines using Legato and group or experiment local servers using the IBM TSM product. TSM is also used for explicit user driven archiving (the pubarch command). This separation is historical and we believe we can now decide on only one of these products. Much of the service is assured by outsourcing and we do not have the time or experience any more to decide how to proceed so have appointed a consultancy to advise us. They are due to report before the end of this year.

Managed Storage

The CASTOR system went into full production at the beginning of this year being successfully used by the second Alice data challenge and replacing HPSS totally for this year's central data recording. There is now some 160 TB of data in 2500000 files in CASTOR as shown in the plot below. This usage has built up faster than we expected and there have been more teething troubles as a result. In addition new functionality was urgently requested by running experiments who were given priority and some of the bugs or features have hence taken longer than we would like to be fixed in the production services. The large amount of data successfully stored strongly demonstrates the overall success of the CASTOR system. In addition we have recruited a new staff member into the team and have put robustness and reliability as the highest priority until we are satisfied, followed by accounting, control and new functionality. In addition work on interfacing Castor to the European Data Grid has started.

We have now put into production a requested wide area ftp interface to CASTOR. Users can make ftp transfers from outside of Cern to the machine wacdr.cern.ch using their CERN afs login and password. They can transfer directly to and from /castor and /afs. Files written into CASTOR will automatically be migrated to tape during the next 30 minutes.

Files to be transferred out of CASTOR will first be staged to the local disks on wacdr so there will be a delay while this happens. For a 1GB file this could be 15 minutes but the staging will complete even if the ftp session is closed. A subsequent ftp request will then find the file to be already on disk and start the transfer.

CERN decided to phase out HPSS in favour of CASTOR. Now that CASTOR phase 1 is in production there is no longer a need for two systems and CASTOR brings us far more flexibility in satisfying the special requirements of HEP, at the price of manpower of course. Owing to workload and vacations we have not systematically started migrating data from HPSS to CASTOR but we are now starting this process. HPSS users should move their files into CASTOR themselves by running the H2C command. When it has successfully completed (which can take many hours for a large volume) users should delete their HPSS files. H2C changes the default file path prefix from HPSS to CASTOR so to delete files use the full path, e.g.,

```
hsm delete hpsssrv1:/hpss/cern.ch/user/m/me/mydir
```

Tapes and Tape Robots

Two years ago we decided on a disaster avoidance scenario of having two physically separate tape robot installations. As a first step we built a second STK robot four-silo complex in the basement of building 513, the first one being on the ground floor. A new building, 613, about 200 metres away is now ready. During the months of October and November we will be moving silos and data from the basement of 513 into the new 613. At the same time we have bought two new silos, each of 6000 cartridge slots, for next year's raw data so as to have sufficient equipment for two complexes of five silos each. We will use the new silos to minimise service downtime in this transition by installing them as the first equipment in the new building. This should be finished by 22 October when we start moving cartridges into them. As soon as an existing silo is empty we will dismantle it and move it

to its new site. The whole process should be finished by the end of November and should be transparent apart from some mount delays except there will probably be two periods of up to twelve hours when a whole complex will be unavailable while new silos are connected up. We will schedule these in quiet times (e.g., accelerator stoppages) with plenty of warning. The CASTOR software itself will be put into a pause mode during these times and batch jobs will be delayed but should not fail.

At the end of 2000 we agreed a project with STK to replace the unreliable helical scan Redwood tape drives with the linear 9940 technology over two years. We now have 28 9940 drives and they have proved very successful. We have already reduced from 32 to 20 Redwood drives and will go down to 12 as part of the move to the new building. By the end of 2002 we plan to have only 4 Redwood drives and for them to be only lightly used. About half of the total number of Redwood cartridges have so far been copied to 9940, mostly as CASTOR files.

An important objective of this timing is to clear the vault to prepare it to be used for the LHC Computing Grid project. Many archive tapes are stored there, plus some DLT tapes, and experiments have already been requested to tell us what we can scrap or to where they should be moved. There will be a limited amount of space in IT for "active" DLTs that can be cycled into the small 600 slot DLT robot. There will be no more manual DLT mounts and we anticipate only small numbers of manual mounts for 4mm and 8mm media.

Also in the area of tape drives we are starting an acquisition process for high quality, high reliability linear technology tape drives to satisfy the requirements of the LHC Computing Grid project for its first phase. This requires a sustained data recording rate to tape of 200MB/sec in 2002 rising to 500MB/sec in 2003. We believe it is too early to rely on the cheaper mid-range technology of LTO for this but will be acquiring a small robot with several LTO vendor tape drives to intensively test this new technology and also support import/export.



4. Desktop Computing

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

4.1 Helpdesk and User Support Contacts

Nicole Crémel and Roger Woolnough, IT / User Services

As explained in several past Computer NewsLetters, last time in CERN-CN-2001-001 (see "Why Helpdesk and User Support are two Entities?", page 1, also at the URL http://ref.cern.ch/CERN/CNL/2001/001/us_mail/), we have tried to disconnect the two e-mail addresses User.Support@cern.ch and Helpdesk@cern.ch (or Computing.Helpdesk@cern.ch). The objective was to make a very clear distinction between the actual "Computing Helpdesk" as it is implemented now, and the "User Support" group in IT division (now renamed to "User Services"). The "Computing Helpdesk" is a front-line (or 1st level) outsourced service where users can raise any type of computing issue via mails, phone calls or visits (a kind of telephone exchange for computing at phone number 78888, with main desk in building 513, and a satellite helpdesk in building 52). "User Support" group, now renamed "User Services" (to eliminate any ambiguity with the Helpdesk), is supervising the helpdesk, and working "behind" it, as a 2nd-level service, to treat a proportion of the queries that are received at the desks.

We have realised that the User.Support@cern.ch e-mail address is so deeply anchored in CERN users' mind and embedded in much documentation that it is inappropriate to completely disconnect this address from the helpdesk.

We have therefore taken "one step back" and re-enabled User.Support@cern.ch, such that users who still use this address (or kept it in their address book) as the entry point for computing issues will be automatically re-directed to the helpdesk. Users will however receive, at the same time, an automatic notification inviting them to use in the future Helpdesk@cern.ch, if they want to raise a computing question. On the other hand, users who wish to contact members of the User Services group in IT division, eventually to report either a good or not-so-good experience with the helpdesk,

or to raise any support issue in general, should use the e-mail address User.Relations@cern.ch.

Following is the text of this notification message, both in English and French:

This is an automatic informative message

(French version is following)

Thanks for mailing User Support. Your message is being delivered to the central Computing Helpdesk where it will be treated or sent to the relevant support-line.

In future please use the address helpdesk@cern.ch for any computing query.

Should you wish to contact members of the User Services Group in IT Division you can use the address User.Relations@cern.ch. Thank You.

Nous avons bien reçu votre message à User.Support@cern.ch qui a été transmis au bureau central d'aide aux utilisateurs pour l'informatique, où il sera traité ou bien dirigé vers la ligne de support appropriée.

La prochaine fois nous vous invitons à envoyer directement toute question informatique à l'adresse électronique helpdesk@cern.ch.

Pour contacter les membres du groupe User Services de la division IT vous pouvez utiliser l'adresse User.Relations@cern.ch. Merci.

About the author(s): *Nicole Crémel and Roger Woolnough are members of the User Services group in IT, and are the CERN responsible for the supervision of operations of the Computing Helpdesk.*

4.2 Reduced Number of Versions of CCDB Tools (userinfo, userreg, spaceadm)

Tony Osborne, IT/CE and Nilo Segura, IT/DB

Given the reality of falling manpower in IT division we are forced to reduce the variety of our tools. userinfo is the CCDB tool which enables users to inspect and partially change personal, email and account information. The client-server versions of userinfo are more powerful in terms of

functionality and speed than the Web version was, and remove the need for specific registration (only about two % or all users have done this so far), which means another password to remember. Therefore the Web version of userinfo is now discontinued. Instead please use one of the following

client-server versions:

- on UNIX: xuserinfo on all supported Unix systems (Lxplus, Rsplus, Dxplus, and Solaris systems such as TH_WGS, NA48_WGS).
- on a PC NICE: From the “CERN phone book”:
 - “Change” button on NICE/95/NT (this version is slightly out of date but we hope it will soon be upgraded)
 - “User Info” button on NICE/2000 (this version is recent and if you do not see this option please reboot your PC)

The latest version of userinfo (not yet available on NICE/95/NT) displays disk space quota/space used for MAILSERV/AFS and gives the possibility to change individual CORE account UNIX shells (for those users liking a complicated life...).

Failing any of the above, either one of your group administrators or, in case you cannot contact any of them, User.Registration@cern.ch can make required changes for you. (You can easily find who are your group administrators in the web phone book form, at the URL <http://consult.cern.ch/xwho>).

Please refer to the URL <http://cern.ch/it/support/userinfo/> for information on the client-server version. Soon the Web URL for userinfo will be changed to point to

the above help page. Actual registrations for using this tool have been removed.

In the same spirit, in order to minimize the number of userreg and spaceadm versions to support, the client-server versions of the tools for group administrators (xuserreg and xspaceadm on Lxplus, Dxplus, Rsplus, Solaris systems) will soon be withdrawn and we will *only* provide the Web version.

Currently the Web version is, we believe, quite acceptable using Internet Explorer (but unacceptably slow with Netscape). To solve this problem we will soon provide userreg and spaceadm versions accessible from an applet viewer from, at least, on LXPLUS.

This means that group administrators should soon register for the Web version of userreg and spaceadm following the instructions given from the starting page of the client-server version of userreg on UNIX systems (xuserreg). If you fail to do this in the next few weeks you will have to request that you be registered for the Web version to User.Registration@cern.ch.

The CCDB tools are not available from HPPLUS since this service is “frozen” at HP/UX 10 and the Oracle software is not compatible with this version of the O/S.

The Web version of userreg is available via the URL <http://userreg.cern.ch>. Likewise, the Web version of spaceadm is available via the URL <http://spaceadm.cern.ch>.



4.3 Computer Security: Prevent Viruses and Worms

Denise Heagerty, IT/IS

Abstract

In July this year several systems at CERN were infected by the “Code Red” worm which resulted in significant network disruption. It is essential that all networked computers remain up-to-date with security patches to reduce the risk of infection from future worms and viruses.

A computer worm is software which copies itself between networked computers, gradually infecting more and more of them. The “Code Red” worm which was launched on the Internet in July this year infected tens of thousands of computers, including 25 on the CERN site. It entered via security holes on computers which had not been regularly patched and seriously disrupted the CERN network with the massive traffic it generated.

System administrators must ensure that their computers have the latest security patches installed, for both the operating system and any application services using the network. If your system is running application services which are not necessary (e.g., some may be running by default or for historical reasons) then disable them. Systems running CERN-certified operating systems can profit from IT Division’s services which apply tested patches automatically or provide

them for manual installation. Those who have installed their own operating systems or applications must themselves check for and install relevant security patches.

Please refer to the CERN’s security recommendations, at URL: <http://cern.ch/security/recommendations> which are reproduced below.

CERN Computer Security Recommendations

Don’t open unexpected e-mail attachments

Viruses often hide in e-mails from strangers, but can also appear to come from someone you know. Opening an attachment can activate a virus and place your computer at risk. If you are not expecting the attachment then either delete the e-mail directly or obtain further details from the sender before opening the attachment. The safest way to read an attachment is to first copy it to disk and then open it using the appropriate program (word, excel, ...). You can also run an anti-virus check on the file before opening it.

Click “cancel” (instead of “ok”) or close unexpected dialogue boxes when using the web.

Visiting a web site sometimes results in dialogue boxes. If you don't expect them or don't understand them then either click “cancel” or close the dialogue box. If you click “ok”, you may be agreeing to transfer and run a file containing a virus.

Don't answer or forward unsolicited e-mail - delete it immediately

We all receive unexpected e-mail: advertising, requests for money or support for a cause. Sometimes it appears to come from an organisation or person that we know, maybe even from someone at CERN. The from address of these e-mails has usually been forged and cannot be trusted. The contents of the e-mail may contain a trick, particularly if it invites you to visit a web site or contains an attachment. If you react to such e-mails you risk introducing a virus into CERN, exposing your personal information (such as your e-mail address and getting even more of these e-mails), and wasting time and money. The more realistic the mail, particularly if it is related to a recent or topical event, the more dangerous it is likely to be. Hoax e-mail warning you of a virus is extremely common—delete it. If the mail asks you to forward it to other people: **Do Not!** Unsolicited e-mail can usually be recognised by checking the subject and sender, so don't even read it—delete it rapidly. If you continue to receive unsolicited e-mail from the same sender then you can report this to abuse@cern.ch.

Run anti-virus software which is automatically updated (several new viruses appear each day).

CERN's centrally managed NICE PCs are equipped with anti-virus software and are automatically updated to limit damage from known viruses. If a virus is discovered, the anti-virus software will notify you, and prevent it from running (by placing it in quarantine). You should continue to work normally, as the anti-virus service will be automatically informed and will contact you if any further action is required. Occasionally, the anti-virus software cannot completely prevent damage, so if you do experience problems contact helpdesk@cern.ch (tel: 78888), with the name of your PC, details of the error message and problem, and request a virus check.

Anyone managing their own Windows PC is responsible for obtaining, installing and keeping their anti-virus software up-to-date. This applies to all PCs on the CERN network, including those of visitors. Regularly updated anti-virus software is particularly important for portable PCs which are used at other locations and connect to other Internet Service Providers since they bypass CERN's security protections. This not only increases their own chance of infection, but places the whole CERN site at risk, since once infected, they can spread an infection from inside our firewall.

Don't copy or run software from non-trusted sources, e.g., via the Internet or physical media such as diskettes or CDs.

Viruses are often hidden inside files. When you copy and run a file containing a virus, you can infect not only your own PC; but can start to spread a virus inside CERN's firewall. Only copy files from trusted sources, such as commercial companies with whom CERN has a software agreement.

Choose secure passwords and change them regularly.

Programs to crack passwords or read them from the network are readily available. To limit the risk of your password being cracked, it should be at least 8 characters long and include letters (both upper and lower case), digits and punctuation. You should change your password regularly and always after a trip where you could have exposed your password at a remote site. More detailed advice is at the URL <http://cern.ch/security/passwords>.

Avoid applications with unencrypted sessions, especially when connecting to CERN from off-site.

Applications such as telnet, ftp and X windows, expose all session data, including passwords, in clear on the network. Using such applications, especially to connect to CERN from other sites, has a strong risk that your password and other personal data will be exposed and used by intruders for malicious activity. You are strongly recommended to use applications, such as *ssh*, which encrypt session data, when accessing CERN from remote sites or for performing sensitive or privileged actions on-site. More detailed advice is at <http://cern.ch/security/ssh>.

Web sites prefixed by “http” expose data in clear. For sensitive data, such as passwords and credit card numbers, ensure that the data is encrypted, e.g., by using web sites prefixed by “https”.

Use CERN's recommended and centrally managed systems

If you manage your own system or have installed your own applications, you are responsible for keeping the software secure:

- limit application services listening on network port numbers to the absolute minimum;
- limit the number of users authorised to access the system to a minimum;
- ensure that the system and applications are securely configured;
- ensure that security patches are regularly applied - this may require upgrading to later versions;
- don't install software which you don't understand;
- respond quickly to fixes proposed by CERN's security team.

Protect your system by CERN's firewall.

Systems connected to CERN's network must be registered at the URL <http://network.cern.ch/register>. The default *OUTGOING* network access allows direct connections to the Internet from CERN, while still offering some protection by CERN's firewall. If your system does not need to access the external Internet and you want extra protection in the firewall, you can register the network access called *NONE*. If absolutely required, *INCOMING* access can be registered, to allow direct access from the Internet. Such systems have significant risk of being attacked and expose the whole CERN

site to security risks. It is essential that they are actively and continually secured (as explained in the previous section).

Keep yourself informed of CERN's security rules and advice:

- *Computing Rules*: <http://cern.ch/ComputingRules>
- *Security Information*: <http://cern.ch/Security>.

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

4.4 Password Recommendations at CERN

Lionel Cons, IT/PDP

Abstract

The password is the most vital part of account security. This article explains why and how to choose good passwords.

Why you need good passwords

The password is the most vital part of account security. If an attacker can discover your password, he/she can use your account to attack systems in or outside CERN, as well as read, modify or delete all your files. CERN's Computing Rules (<http://www.cern.ch/ComputingRules>) require that you protect your accounts with a good password:

III.11: All accounts must have appropriate access protection, such as account codes or passwords.

III.12: The user shall take the necessary precautions to protect his personal computer or work station against unauthorized access. The user shall also protect details of his personal account, particularly by avoiding obvious passwords and shall not divulge his passwords to any third party, unless expressly authorized by his Division Leader. Upon request from the CERN Computer Security Officer or the service manager concerned, the user shall select a new password.

How to choose good passwords

A good password is:

- *private*: it is used and known by one person only;
- *secret*: it does not appear in clear text in any file or program or on a piece of paper pinned to the terminal;
- *easily remembered*: so there is no need to write it down;
- *not guessable by any program* in a reasonable time, for instance less than one week;
- *longer than 7 characters*;
- *a mixture of upper/lower case, letters, digits and punctuation*.

Here are some *hints* to help you choose good passwords:

- Choose a line or two from a song or poem, and use the first letter of each word. For example,
 "In Xanadu did Kubla Kahn a stately pleasure dome decree"
 becomes *"IXdKKaspdd"*.
- Alternate between one consonant and one or two vowels with mixed upper/lower case. This provides non-sense words that are usually pronounceable, and thus easily remembered. For example: *"roUtboo"* or *"quADpop"*.
- Choose two short words (or a big one that you split) and concatenate them together with one or more punctuation characters between them, e.g., *"dog+F18"* or *"comP!lUTer"*.

Attackers and programs that can try to break in to your account know a large number of "frequently used" passwords. Here are some *guidelines to avoid "guessable" passwords*:

- *don't* use your login name in any form (as-is, reversed, capitalised, doubled, with a prefix, with a suffix, etc.).
- *don't* use in any form your first or last name and, more generally, any information easily obtained about you. This includes car license plate numbers, telephone numbers, insurance numbers, the brand of your car, the name of the street you live on, the name of your spouse or of your children, etc.
- *don't* use a word contained in any dictionary of any language, spelling lists, or other lists of words (acronyms, sequences of letters like *"abcdef"* or *"qwerty"*, place names, car names, cartoon heroes...).

Why you must change passwords

Even if you choose a good password, it can still be discovered: someone may see you typing it or capture it by snooping on the computer or network. If you accidentally type

your password in place your login name, it may appear in system log files:

```
joe      tty9      Wed Apr 28 09:37
XSecret! pty/ttys0 Fri Feb 26 15:15 - 15:16 (00:00)
fred     pty/ttys0 Fri Feb 26 15:16 - 14:27 (87+22:11)
```

For all these reasons you must change your passwords from time to time. We recommend that you *change your passwords whenever you return from a trip* that could have exposed them and with a *minimum frequency of twice a year*.

References

For more information, you can consult:

- Password Crackers - Ensuring the Security of Your Password (<http://www.securityfocus.com/infocus/1192>)
- Password Primer from the University of North Carolina at Chapel Hill. (<http://www.cs.unc.edu/cgi-bin/howto?howto=security-passwords>)
- SLAC Account Password Policy (<http://www.slac.stanford.edu/comp/security/passpolicy.html>)

About the author(s): *Lionel Cons is a member of the CERN security team (Computer.Security@cern.ch).*



4.5 HEPiX News

Maria Dimou, IT/User Services

Abstract

HEPiX (as described by its chairperson, Lisa Giacchetti) is an international group of UNIX users and system administrators working in cooperating High Energy Physics institutions.

The focus of the group is to share experiences, influence vendors and standards bodies and to investigate solutions to those problems which impede the use of Unix for mainstream work at the member institutions. Currently less emphasis is being placed on influencing outside bodies and more on the sharing of experiences, code and documentation.

The group meets twice per year, in spring at the premises of a European HEP institute, and in the autumn in a USA site.

Every site submits a number of talks and a report with the latest news of main computing services. CERN's site report at the last meeting, that took place in Oakland, California on October 15-18, 2001, can be found at the URL <http://cern.ch/ref/hepixon/2001/005>. More generally, the URL <http://cern.ch/ref/hepixon/> contains recent reports and links to other HEPiX web sites.

About the author(s): *Maria Dimou is the CERN representative in HEPiX.*



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 Upgrade of the NICEWWW Service

Per Hagen, IT/IS

NICEWWW, the service which implements web access to the NICE home directories, has been upgraded last September 2001. The new computer WEB10 has taken over the service from WEB2.

A new version of the NICEWWW gateway software (ISAPI filter) has replaced the current one. The following issues have been addressed in the ISAPI filter:

File handle leak

The ISAPI filter introduced during the summer contains a file handle leak such that the server will run out of non-pageable memory after a few weeks, depending upon load. This problem has been fixed.

Default page

The ISAPI filter now handles default pages as for the central server. Previously, only `default.htm` was recognized as a default page. Now, `default.html`, `index.htm*` and `welcome.*` are also supported. The reasons for supporting all these is for historical consistency between Windows and UNIX web servers.

Cache coherence

The NICEWWW computer is using a disk cache (> 30 GB) in order to efficiently serve pages between the web browser and the home directory servers.

The biggest issue since the birth of the NICEWWW software is that while pages were correctly brought into the cache "just in time", there was no proper handling to delete cache entries where the original had been deleted by the author.

This was a "nice" feature in the sense that when a home server was down, the NICEWWW web server could happily continue serving pages. To compensate for this artifact a scheduled job would delete pages which were not referenced for more than a few days (based on the last accessed date/time). This algorithm left 2 major problems unsolved:

- The time between a user deleting web contents and the web server deleting the cache entry was of the order of days
- A popular page might never expire in spite of being deleted. In particular world wide search engines would reread the page regularly.

During the summer the ISAPI filter was updated to delete cache entries "just in time" when discovering that the original contents no longer exist. The role of the scheduled job

has changed to trim the overall size of the cache. This allows for using a larger cache size. The new server has currently 36 GB available for cache + log files.

This new software release changes the following behaviour: when a home server is down the corresponding cache is allowed to continue according to the uncertainty principle that the probability the page should be served is on average high. The caching algorithm would fail if there was a type mismatch between cache and original.

Scenario: User creates a page named X. X is read via the web and ends up in the cache. User deletes X and creates a directory called X. Inside the directory X the user creates the page Y. The page Y cannot be read via the web because the X in the cache is a file...

Performance optimized for Windows/2000 + IIS5

For historical reason the ISAPI filter was changing security context for each request. This is not necessary as the IIS web server is already running with the correct security context and furthermore only anonymous browse access is supported such that the security context is invariant for the lifetime of the web server.

For historical reasons the ISAPI filter was issuing network connect and disconnect requests for each request. This is for example not needed for Windows/2000 and DFS which is handled transparently by the lower layers in the file system.

Use of non-registered site names

The ISAPI filter contained a historical feature which allowed the use of non-registered site names. Prior to spring 2000 NICEWWW web sites did not have to be registered and the syntax `~username` was the way to access the service. The change in the spring 2000 mandate that all sites must be centrally registered and moderated. In addition, the `~` character is deprecated for site names as it cannot be used in the new web naming scheme.

The worst case scenario is that using `~sitenames` could enable a user to get access to web contents where the owner had not intended web access although the contents would need to be public readable (*Everyone* or *Domain Users*).

This release suppresses this behaviour to enforce conformance to web policy and privacy.

Directory browsing no longer returns link to web root

The top-level page of the directory browsing used to return a link to the root of the web server. This is outside the site and consequently the NICEWWW web server would redirect the web browser to the default CNDCI site. This was rather confusing for the user and not useful and the link has therefore been suppressed.

Quality assurance

A new approach for debugging all of the code interactively has been used. A mechanism for real time diagnostics has been retained.

Please address any question about this change directly to the author (Per Hagen, IT/IS).

5.2 The CERN Search Engine

Dmitry S. Nikitov and Per Hagen, IT/IS

The *CERN Search Engine* is the service that allows users to search the huge amount of information stored inside the CERN Intranet.

The CERN Intranet today consists of roughly 700 000 documents generated by numerous applications. The *Search Engine* is able to index and search documents in popular formats: HTML, XML, PDF, Microsoft Office documents (PowerPoint, Word). Indexing is independent of platform because the server is using the standard Internet protocols.

In addition to the collection of CERN documents in the Intranet, the *Search Engine* is providing full-text search of the impressive collection of scientific documents in the CERN Document Server. Furthermore, a few experiments are using the *service* for private information.

The *Search Engine* serves today over 1000 queries per day and more than 200 queries/hour during working hours. You can use either the *Simple* or the *Advanced Search* forms for complex criteria, like "date"-searches, without having to learn the search syntax.

The *Search Engine* has been serving CERN for several years and showed itself as a reliable and stable service. However, as time is passing, software and hardware needs some maintenance to make it up-to-date.

The *CERN Search Engine* is based upon *Inktomi Enterprise Search*. The company name changed from *Infoseek* to *Inktomi* in 2000.

What has been done

We have tested and configured version 4.2.

The current version we are using is 4.0.8. The changes in the new version 4.2 are mainly bug-fixes, security-fixes and of cosmetic nature.

We have taken the opportunity to make design changes to the search page and update the links

The new service will run on the well-known port 80 while the old service is running on port 8765. Using the standard port 80 solves the problem that some external sites cannot reach the search service due to firewalls. It is planned to redirect the old port 8765 after the old service ceases to exist such that the old URL is still operational.

We have dropped the support for *Topics* on the search page in the new service. We have instead provided a link to the CERN Index on the CERN Home page.

What is next

The current service will continue to run during a transition period towards the end of the year.

Web site owners registered in our central database will be notified by email about the new service. A notice about the change, and a link to the new service, will be added to the search page.

Web authors are encouraged to update pages to reflect the use of the new URL.

The current service is located at the URL <http://wwwsearch:8765/>, while the new service is located at the URL <http://cern.ch/search>.

Some search examples

The following short examples should give the reader some idea how to search. See the online help for complete description.

The search syntax is Inktomi's Natural Language.

Syntax **Explanation and example**

word Search for documents containing "word".

Example: *higgs*

"word1 word2"

Search for a *term* of adjacent words.

Example: "higgs boson"

Albert Einstein: Use uppercase (capital) letters to identify proper names. This is equivalent to the term "Albert Einstein"

W3C

Uppercase matches only exact case.

That is, W3C does not match w3C w3c W3c. lowercase matches both cases.

term1 term2 this *or* that.

Search must satisfy at least one criterium.

Example: "top quark" "beauty quark"

+term this *and* that. Search must satisfy several criteria.
Example: +lepton +"Higgs boson"
-term should *not*. Try to exclude given terms.
Example: +"proton decay" -SUSY -"super
symmetry"

Upgrade Plans

All details concerning the upgrade of the central CERN search engine are available at the URL: <http://cern.ch/WebServices/Operations/01112001.htm>.

5.3 Network Portable Outlets for DHCP Service Users

Jean-Michel Jouanigot, IT/CS

Abstract

Change of IP addresses on Network Portable Outlets.

In order to improve the quality of the DHCP service on the structured cabling portable outlets, and increase the number of addresses we can provide to the DHCP users, the setup of all Portable Outlets is progressively being changed.

The only visible impact for laptops (portable PCs and MACs) making use of DHCP on portable outlets is that the dynamically allocated addresses will become of the form 128.141.x.x instead of 137.138.x.x. (Some parts of the old 128.141 network are recycled for addressing with DHCP). Your system parameters will automatically be set to the correct values (default gateway, etc.), and you should not alter them in any way. Any optional fixed address/addresses (indicated on a sheet which should normally be on display close to the outlet) will stay of the form 137.138.x.x.

In addition, in order to prevent incorrectly configured private BOOTP/DHCP servers from disrupting the DHCP service, all DHCP/BOOTP requests from a system connected to a

portable outlet will be directed to the Central BOOTP/DHCP server managed by IT. As a consequence, booting diskless systems such as VME boards (using BOOTP/DHCP) connected to a portable plug will **not** be supported, even if you run your own BOOTP/DHCP server. Other network broadcast packets (e.g., SUN RPC yellow page service) will also be blocked. Finally, systems acting as BOOTP/DHCP servers must **not** be connected to portable outlets.

At some future date, for network security reasons, only correctly registered computers will be allocated a DHCP address permitting the full range of access possibilities to the general CERN network. This restriction will be based on the computer hardware address (also known as MAC address, adapter address or ethernet address), so we strongly encourage all laptop users to check or register their systems at the URL <http://network.cern.ch/register>.

General help for the installation and use of portable computers may be found at the URL:

<http://cern.ch/it-div-cs/public/services/notebooks.html>.

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 Engineering Applications on Windows 2000, Availability Planning

John Evans, IT/CE

The availability for all engineering programs installed by IT/CE group under NICE 2000 can be found at the URL:

<http://wwwinfo.cern.ch/ce/pc/w2k-status.html>

Please make sure that all programs you need are available before you migrate to NICE 2000.

The status at the end of November is detailed below, but refer to the latest version of that web page to see the latest status. In each case we show the name of the application, its manufacturer (in *italic*) and its availability.

Electronics

Inspect 1.1 *CAE Consulting* is available.

PCAD 2001 *Altium* is available.

PSpice *Cadence* is available.

Simplorer 4.2 *Ansoft* is available.

Timing Designer *Cronology* will be available mid-December 2001.

Mathematics

Axum 6 *Mathsoft* is available.

MathCAD 2001 *Mathsoft* is available.

Mathematica 4.1 *Wolfram* is available.

Matlab 6 R12 *Mathworks* is available.

VisSim 4.5 *Mathsoft* will be available mid-December 2001.

Mechanics

AutoCAD 2002 *AutoDesk* is available.

Structural Analysis and Field Calculation

Ansys 5.7 *Ansys* is available.

Cetim 2.03 *Cetim* is available.

Design Space *Ansys* will be available mid-December 2001.

Esacomp 2.0 *Componeering* will be available mid-December 2001.

MicroWave Studio 3 *CST* is available.

Superfish 6.03 *Los Alamos Labs* is available.

Other Tools

Cadim 2.3 *Eigner + Partner* is available.

CEHSM *CERN* is available.

CESTAT *CERN* is available.

HPGL View 5.01 *CERN* is available.

LabVIEW Please see <http://pcitco25.cern.ch/SI/lvbw/>.

See 3000 *IGE+XAO* will be available mid-December 2001.

Tracelec *TRACELEC* will be available mid-December 2001.

Visio 2002 is provided by IT/IS and is available.

VoloView Express 1.14 *Autodesk* is available.

WorkView 1.2.C *Delta Concept* is available.

The dates for not yet available applications assume that there are no problems with software availability, or compatibility issues with CERN Win2k systems.



6.2 Monte Carlo Library News

Ian McLaren, IT/API

New versions of some Fortran based Monte Carlo libraries have become available since the summer. They will be installed in a Cernlib 2002 ("new") pre-release in a few weeks. For detailed information, links to the relevant WEB pages are given below. As most users know there are often dependencies on specific versions of other libraries, so the latest manuals and release notes should be consulted.

Ariadne 4.11

Ariadne can now be used together with the PYTHIA program of version 6.1 or higher as well as with older versions of JETSET (7.4) and PYTHIA (5.6), but separate libraries have to be used. All internal floating point variables in ARIADNE are now in double precision. The parameters PARA in /ARDAT1/ are still in single precision, though. There is an updated manual, examples and down-loads at the URL <http://www.thep.lu.se/~leif/ariadne/>.

Pythia 6.2

Pythia 6.2 is now available on the standard web page,

<http://www.thep.lu.se/~torbjorn/Pythia.html>

The main improvement is that Pythia 6.2 comes with a completely updated big manual (around 420 pages). The possibility to have a code and documentation in phase is a good reason to recommend all experimental collaborations to upgrade to this new version. There are a few points of backwards incompatibility, described in the update notes, but an upgrade from Pythia 6.1 should not require a major effort.

Isajet 7.58

This new version (replacing version 7.51) contains a number of improvements. The Cernlib version will be compiled with the RANLUX option. Full details and an updated manual are at the URL

<http://www.phy.bnl.gov/~isajet>

stdhep v4_09

As with previous versions, the libraries, executables and include files for the FNAL supported platforms will be available also from the Cernlib release. Full details are at the URL

<http://www-pat.fnal.gov/stdhep>

6.3 Software Development Tools Service

Eric Poinsignon, IT/PS

Abstract

Description, content and news from the SDT service.

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

Latest news

- *Rational Unify Process 2001A.04*: (RUP) is available. It's a Web-enabled software engineering process for software best practices. We invite you to *try it!* RUP is a kind of web encyclopedia for *any software-related work*. Have a 15 minutes tutorial tour starting from the URL
<http://s.home.cern.ch/s/swtool/www/pttools/RUP/rup/RationalUnifiedProcess/index.htm>
You can an SDT slide presentation at the URL
http://sdt.cern.ch/RUP/RUPQuickTour_files/frame.htm
More details in the RUP home page:
<http://sdt.cern.ch/RUP/>
- *Together 5.0* (<http://sdt.cern.ch/Together/>): a UML Integrated Development Environment generating Java and C++. The Nodelock academic licenses are now available. This allows all *local installations* to work.

- *Rose 2001* (<http://sdt.cern.ch/Rose/ProcLinux.html>), an OMT/Booch/UML analysis and design tool with C++/Java/Ada code generation is now available also on Linux.
- *CodeWizard 4.0* (<http://sdt.cern.ch/CodeWizard/>), coding standards enforcement tool that prevents errors and standardizes C/C++ code is in production.
- *Soda for Frame* (<http://sdt.cern.ch/Soda/>), a software documentation automation tool is no longer used, so that the maintenance contract has been stopped. Soda for Word will continue.
- *SDLT templates for FrameMaker 6.0* (<http://docsys.web.cern.ch/docsys/framemaker/templates/sdlt/>), the Software Documentation Layout Templates) has been upgraded to v2.1. The template has no new features, and it is usable without any upgrade.
- *CERN-Specific Office templates for Frame* (<http://docsys.web.cern.ch/docsys/framemaker/templates/cernoffice.html>) have been upgraded for FM 6.0. FM users on Unix will get them directly in the FM "New" menu. Other users will find at the above URL.

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

7. Desktop Publishing

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

7.1 Installing $\text{\LaTeX}$ 2001

Michel Goossens, IT/API

The $\text{\TeX}$ system installed at CERN is the $\text{\TeX}$ live 4 distribution (June 1999). Since this release is well over two years old I have installed $\text{\TeX}$ live 7 (November 2001) on AFS.

To have access to this release, on Linux add the directory /afs/cern.ch/sw/ $\text{\TeX}$ live/tl7/bin/i386-linux and on Solaris the directory below to your PATH variable: /afs/cern.ch/sw/ $\text{\TeX}$ live/tl7/bin/sparc-solaris2.7

For this release $\text{\LaTeX}$ is compiled in its “huge” variant and comes with the hyphenation patterns of over twenty languages preloaded, as shown below.

```
> PATH=/afs/cern.ch/sw/ $\text{\TeX}$ live/6alpha/bin/i386-linux:$PATH
> export PATH
> latex
This is TeX, Version 3.14159 (Web2C 7.3.3.1)
**\ relax
LaTeX2e <2001/06/01>
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.
```

On NICE2000 this release is available in the dfs directory \\cern.ch\dfs\Experiments\sw\textlive in the folder

tl7. If you map the above dfs directory to your “T drive”, then to run the $\text{\TeX}$ -related programs on NICE 2000 you should add the binary directory to your Path environment variable using the following procedure:

```
Settings
- Control Panel
- Double click on System
- Choose "Advanced" tab
- Click on "Environment Variables..."
- Go into lower frame
- Move down until you see "Path" variable
- select it and click on "Edit..."
- add T:\tl7\bin\win32; to end of path and OK
```

Now you should be able to use latex, dvips, etc. directly (they will automatically be found on the “T drive”). Early next year I plan to make available a procedure to install the files locally on your hard disk, together with other interesting $\text{\TeX}$ -related tools.

Note that new versions of $\text{\TeX}$ are only made available on Linux, Solaris, and NICE2000. The other Unix (and Windows platforms) being “frozen” and will *no longer be upgraded*. Hence they will stay with the release of June 1999.



7.2 XML Applications at CERN

Michel Goossens, IT/API

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

The XML tools, such as saxon.sh, xerces.sh, batik.sh, can be executed by defining the shell variable XMLCDROM and by including the relevant directory in your PATH variable, as follows (depending on your Unix shell):

```
# for Bourne-like shells
XMLCDROM=/afs/cern.ch/sw/XML/cdrom
PATH=$XMLCDROM/bin/i386:$PATH
export PATH XMLCDROM
```

```
# for C-like shells
setenv XMLCDROM /afs/cern.ch/sw/XML/cdrom
setenv PATH $XMLCDROM/bin/i386:${PATH}
```

I am in the process of making all the XML-related files just mentioned also available on the NICE2000 system at CERN. At present a few of the tools (and an earlier version of the doc and tutorials) are already installed in the dfs directory \\cern.ch\dfs\Experiments\sw\Xml. If you map this folder, e.g., to your “X drive”, and add the folder X:\bin\win32; to the end of your path, as explained in the previous section, you can use saxon, etc. For the moment on NICE2000 I have hard-code the path using the “X” drive into the bat files. I intend to make these files more independent of the precise path where the tools are installed. Moreover, as in the case of $\text{\TeX}$, I am looking into ways of easing the installation of the tools locally.

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 A Quick Tour of Compiling, Linking, Loading, and Handling Libraries on Unix

Michel Goossens and Ian McLaren, IT/API

Abstract

In the September/October 2001 issue of Computing in Science & Engineering, an I.E.E.E. publication, edited by Paul F. Dubois, three authors, David M. Beazley, Brian D. Ward, and Ian R. Cooke, wrote an article “The Inside Story on Shared Libraries and Dynamic Loading”. As we thought that the subject was also interesting for the readers of the CNL, we obtained permission of the I.E.E.E. to make available the original version as a PDF reprint, see

<http://ojps.aip.org/journal.cgi/dbt?KEY=CSENF&Volume=CURVOL&Issue=CURISS#MINOR5>

with the following Copyright:

© 2001 IEEE. Reprinted, with permission, from Computing in Science & Engineering–September 2001 Volume 3, Issue 5.

As a complement to that article we present here a few practical examples on how to handle libraries with C++ programs, e.g., NAGLIB, CLHEP, etc., that are used at CERN.

Compilers and object files

When building an executable program the compiler converts source files into object files. Each object file contains machine code instructions that correspond to statements in the program source. Generally, object files are broken into a collection of sections corresponding to different parts of the source program. As an example, let us consider the following small C++ program. It is good practice to explicitly specify the namespace `std::` when referring to *standard* classes and methods. This will guarantee maximal compatibility with standard-compliant C++ compilers and run-time libraries.

```
#include <iostream>
int main(int, char **) {
    int ny = 2002;
    std::cout << "Merry Xmas, Happy " << ny << "!" << std::endl;
}
```

We can compile it with a C++ compiler (the `g++` GNU compiler on Linux, and the native `c++` compiler on Solaris, but whatever is explained in the present article should apply to other standard-compliant compilers on different computer platforms as well), and then execute the generated executable program, first on Linux (e.g., `lxplus` at CERN):

```
$ g++ merry.cpp -omerry
$ merry
Merry Xmas, Happy 2002!
```

or similarly on Sun OS 5.7 (e.g., `sundev` at CERN):

```
$ CC merry.cpp -o merry1
$ merry1
Merry Xmas, Happy 2002!
```

It comes as no surprise that the generated output is identical on both systems. Information about the object file that contains the various compiled sections (machine code instructions, global and local data) and the symbol table describing the source program identifiers can be obtained by the Unix `nm`, command, which lists symbols from object files. The detailed information provided by the command `nm` is platform-dependent, and the man pages should be consulted to understand its precise structure. However, we show some salient features below, first on Linux:

```
$ nm --print-file-name --demangle merry
merry:080499ac A _DYNAMIC
merry:08049978 A _GLOBAL_OFFSET_TABLE_
```

```

...
merry:08049a60 B cout
merry:08049928 W data_start
merry:      U endl(ostream &)
merry:08048834 t fini_dummy
merry:08049934 d force_to_data
merry:08049934 d force_to_data
merry:0804883c t frame_dummy
...
merry:08048864 T main
merry:08049a8c b object.8
merry:0804992c d p.2

```

and then on SunOS 5.7 (more than 2000 lines).

```

$ nm -V -A -R -C merry1
nm: Software Generation Utilities (SGU) Solaris-ELF (4.0)

```

```
merry1:
```

merry1: [Index]	Value	Size	Type	Bind	Other	Shndx	Name
merry1: [25]		0	0 SECT	LOCL	0	24	merry1:
...							
merry1: [1274]		372276	0 FUNC	GLOB	0	UNDEF	merry1:.div
merry1: [790]		372264	0 FUNC	GLOB	0	UNDEF	merry1:.mul
merry1: [582]		372072	0 FUNC	GLOB	0	UNDEF	merry1:.stret8
...							
merry1: [389]		0	0 FILE	LOCL	0	ABS	merry1:xsZ9_GTZj8hTFL_518EJ
merry1: [374]		0	0 FILE	LOCL	0	ABS	merry1:yFOC_cVsuFXDicukF3CL
merry1: [388]		0	0 FILE	LOCL	0	ABS	merry1:yJhYas4P93YJaULD27v0
merry1: [162]		0	0 FILE	LOCL	0	ABS	merry1:z0tiZS4hQ7CvMDXtQoDv

Linkers and linking

To build an executable file, the linker (for example, `ld`) collects object files and libraries. The linker's primary function is to bind symbolic names to memory addresses. To do this, it first scans the object files and concatenates the object file sections to form one large file (text sections of all object files are concatenated, data sections are concatenated, and so on). Then, it makes a second pass on the resulting file to bind symbol names to real memory addresses. To complete the second pass, each object file contains a relocation list, which contains symbol names and offsets within the object file that must be patched. For example, the relocation list for the object file for our simple C++ example looks something like this:

```

$ objdump -R merry

merry:      file format elf32-i386

DYNAMIC RELOCATION RECORDS
OFFSET      TYPE          VALUE
080499a8 R_386_GLOB_DAT  __gmon_start__
08049a60 R_386_COPY      cout
08049984 R_386_JUMP_SLOT endl__FR7ostream
08049988 R_386_JUMP_SLOT __ls__7ostreamPFR7ostreamR7ostream
0804998c R_386_JUMP_SLOT __throw
08049990 R_386_JUMP_SLOT __ls__7ostreamPCc
08049994 R_386_JUMP_SLOT __deregister_frame_info
08049998 R_386_JUMP_SLOT __libc_start_main
0804999c R_386_JUMP_SLOT __ls__7ostreami
080499a0 R_386_JUMP_SLOT __register_frame_info
080499a4 R_386_JUMP_SLOT __gmon_start__

```

Similarly, on SunOS 5.7 the `dump` command with the `-L` option shows information about dynamic linking and static shared libraries for an object file.

Static and shared libraries

To improve modularity and reusability one often groups commonly-used functions in *libraries*. In the past, such a library was in the form of an archive (`.a` file), created like this:

```

$ ar cr libmyar.a myobj1.o myobj2.o myobj3.o ...

```

The resulting `libmyar.a` file is called a *static* library. Such an archive is merely a collection of raw object files strung together along with a table of contents for fast symbol access (sometimes one even had to manually construct this table of contents with a special utility, such as `ranlib`.) Static libraries are easy to create and use, but they present a number of software maintenance and resource utilization problems. Hence, nowadays shared or dynamic link libraries are preferred.

When a static library is included during program linking, the linker makes a pass through the library and *adds all the code and data* corresponding to symbols used in the source program. Thus, when the linker includes a static library in a program, it copies data from the library to the target program. This means that whenever a change to a library is needed, everything linked against that library must be rebuilt for the changes to propagate. Moreover, copying the contents of a library into the target program wastes disk space and memory. For example, if every executable on a system includes its own copy of the C library, the size of these programs increases dramatically, and when active, they waste a considerable amount of system memory because they each store their copies of the library functions.

It is important to note that the linker ignores unreferenced library symbols and aborts with an error when it encounters a redefined symbol.

These days the maintenance and resource problems of static libraries are addressed by using shared libraries or dynamic link libraries (DLLs). The primary difference between static and shared libraries is that using shared libraries delays the actual task of linking to runtime, where it is performed by a special dynamic linker-loader. So, a program and its libraries remain decoupled until the program actually runs. This means that when a bug has to be corrected in a commonly used library the latter can be patched and updated without the need to re-compile and re-link applications referring to the library in question, whose code is only dereferenced when the program is actually executed. Moreover, optimizations on the system level are possible, e.g., by grouping non self-modifying executable instructions in read-only memory regions that can be shared among processes. Thus, if many programs are running and a large fraction of them include the same library (e.g., the C library), the operating system can load a single shared copy of the library's instructions into physical memory. This significantly reduces memory use and improves system performance.

On most systems, the static linker handles both static and shared libraries. In fact, `gcc` and `g++` are set up to use shared libraries by default. In the example below we first compile and link in static mode, and then in (default) shared mode. In the latter case the location of the shared libraries is determined by the value of the `-L` option of the `g++` command at loading and the setting of the environment variable `LD_LIBRARY_PATH` at execution, as explained later. Library dependencies are displayed by the `ldd` (list dynamic dependencies of executable files or shared objects) command.

```
# Generate static version of program
# add -v option to g++ command to see how libraries are searched
$ g++ --static merry.cpp -o merryst
$ ls -l merryst
-rwxrwxr-x 1 goossens dl      1436909 Nov 14 15:51 merryst
$ g++ merry.cpp -lpthread -o merrysh
$ ls -l merrysh
-rwxrwxr-x 1 goossens dl      13341 Nov 14 16:03 merrysh
$ ldd merrysh
      libpthread.so.0 => /lib/libpthread.so.0 (0x2aacb000)
      libstdc++-libc6.1-1.so.2 => /usr/lib/libstdc++-libc6.1-1.so.2 (0x2aade000)
      libm.so.6 => /lib/libm.so.6 (0x2ab20000)
      libc.so.6 => /lib/libc.so.6 (0x2ab3d000)
      /lib/ld-linux.so.2 => /lib/ld-linux.so.2 (0x2aaab000)
```

Similarly, on SunOS 5.7 this becomes:

```
# Generate static version of program
$ CC -Bstatic merry.cpp -o merryst1
$ ls -l merryst1
-rwxrwxr-x 1 goossens dl      789912 Nov 15 15:55 merryst1
$ CC merry.cpp -lpthread -o merrysh1
$ ls -l merrysh1
-rwxrwxr-x 1 goossens dl      338240 Nov 15 15:56 merrysh1
$ ldd merrysh1
      /usr/platform/SUNW,Ultra-60/lib/libc_psr.so.1
      libpthread.so.1 => /usr/lib/libpthread.so.1
      libCrun.so.1 => /usr/lib/libCrun.so.1
      libm.so.1 => /opt/SUNWsp/1/lib/libm.so.1
      libw.so.1 => /usr/lib/libw.so.1
      libc.so.1 => /usr/lib/libc.so.1
      libdl.so.1 => /usr/lib/libdl.so.1
      libthread.so.1 => /usr/lib/libthread.so.1
      /usr/platform/SUNW,Ultra-60/lib/libc_psr.so.1
```

When binding symbols at runtime, the dynamic linker searches libraries in the same order as they were specified on the link line and uses the first definition of the symbol encountered. If more than one library happens to define the same symbol, only the first definition applies. Although symbol binding is normally of not much interest to the final user, sometimes useful information can be obtained by studying the (huge) output generated by setting the LD_DEBUG environment variable to the value bindings before starting your program. For instance, the following command gives almost 1000 lines of output: (we only show the first two and last nine...).

```
$ LD_DEBUG=bindings merrysh
27534: binding file /lib/libc.so.6 to /lib/libc.so.6: symbol '_nl_current_LC_CTYPE' [GLIBC_2.0]
27534: binding file /lib/libc.so.6 to /lib/libc.so.6: symbol '_nl_current_LC_COLLATE' [GLIBC_2.0]
27534: ...
27534: calling fini: /lib/libm.so.6
27534:
27534: binding file /lib/libm.so.6 to /usr/lib/libstdc++-libc6.1-1.so.2: symbol '__deregister_frame_info' [GLIBC_2.0]
27534:
27534: calling fini: /lib/libc.so.6
27534:
27534: binding file /lib/libc.so.6 to /usr/lib/libstdc++-libc6.1-1.so.2: symbol '__deregister_frame_info' [GLIBC_2.0]
27534: binding file /usr/lib/libstdc++-libc6.1-1.so.2 to /lib/libc.so.6: symbol '_IO_file_setbuf' [GLIBC_2.1]
27534: binding file /lib/libc.so.6 to /lib/libc.so.6: symbol '_exit' [GLIBC_2.0]
```

Library loading

When a program linked with shared libraries runs, program execution does not immediately start with that program's first statement. Instead, the operating system loads and executes the dynamic linker (usually called `ld.so`), which then scans the list of library names embedded in the executable. These library names are never encoded with absolute pathnames. Instead, the list has only simple names such as `libpthread.so.0`, `libm.so.6`, and `libc.so.6`, where the last digit is a library version number. To locate the libraries, the dynamic linker uses a configurable library search path. This path's default value is normally stored in a system configuration file such as `/etc/ld.so.conf`, which on `lxplus` at CERN contains the following directories:

```
/lib
/usr/lib
/usr/X11R6/lib
/usr/i486-linux-libc5/lib
/usr/local/lib
```

If one wants to extend the set of directories searched one should add the paths to these directories to the LD_LIBRARY_PATH environment variable. It should be mentioned that libraries load in the order in which they were linked, thus for the following command:

```
$ g++ mysource.cc -lmy1 -lmy2 -lmy3 ...
```

the dynamic linker loads the libraries in the order `libmy1.so`, `libmy2.so`, `libmy3.so`, etc. If a shared library includes additional library dependencies, those libraries are appended to the end of the library list during loading. For example, if the library `libmy2.so` depends on an additional library `libmy4.so`, that library loads after all the other libraries on the link line (in our example after `libmy3.so`). The load process guarantees that no library is ever loaded more than once with previously loaded copies being used when a library is repeated.

You can follow in detail how the dynamic linker loads libraries by setting the LD_DEBUG environment variable to `libs`, e.g.,

```
$ LD_DEBUG=libs merry
14224: find library=libstdc++-libc6.1-1.so.2; searching
14224: search cache=/etc/ld.so.cache
14224: trying file=/usr/lib/libstdc++-libc6.1-1.so.2
14224:
14224: find library=libm.so.6; searching
14224: search cache=/etc/ld.so.cache
14224: trying file=/lib/libm.so.6
14224:
14224: find library=libc.so.6; searching
14224: search cache=/etc/ld.so.cache
14224: trying file=/lib/libc.so.6
14224:
14224: calling init: /lib/libc.so.6
14224:
14224: calling init: /lib/libm.so.6
```

```
14224:
14224: calling init: /usr/lib/libstdc++-libc6.1-1.so.2
14224:
14224: initialize libc
14224:
14224: initialize program: merry
14224:
14224: transferring control: merry
14224:
Merry Xmas, Happy 2002!
14224:
14224: calling fini: /usr/lib/libstdc++-libc6.1-1.so.2
14224:
14224: calling fini: /lib/libm.so.6
14224:
14224: calling fini: /lib/libc.so.6
```

and on SunOS 5.7

```
$ LD_DEBUG=libs merry1
```

```
28947:
28947: configuration file=/var/ld/ld.config: unable to process file
28947:
28947: find library=libcun.so.1; searching
28947: search path=/opt/SUNWspro/lib:/usr/openwin/lib:/usr/dt/lib:/usr/local/lib (LD_LIBRARY_PATH)
28947: trying path=/opt/SUNWspro/lib/libcun.so.1
28947: trying path=/usr/openwin/lib/libcun.so.1
28947: trying path=/usr/dt/lib/libcun.so.1
28947: trying path=/usr/local/lib/libcun.so.1
28947: search path=/opt/SUNWspro/lib/rw7:/opt/SUNWspro/lib:/opt/SUNWspro/lib:/usr/ccs/lib:/usr/lib (RPATH from file merry1)
28947: trying path=/opt/SUNWspro/lib/rw7/libcun.so.1
28947: trying path=/opt/SUNWspro/lib/libcun.so.1
28947: trying path=/opt/SUNWspro/lib/libcun.so.1
28947: trying path=/usr/ccs/lib/libcun.so.1
28947: trying path=/usr/lib/libcun.so.1
28947:
28947: find library=libm.so.1; searching
28947: search path=/opt/SUNWspro/lib:/usr/openwin/lib:/usr/dt/lib:/usr/local/lib (LD_LIBRARY_PATH)
28947: trying path=/opt/SUNWspro/lib/libm.so.1
28947:
28947: find library=libw.so.1; searching
28947: search path=/opt/SUNWspro/lib:/usr/openwin/lib:/usr/dt/lib:/usr/local/lib (LD_LIBRARY_PATH)
28947: trying path=/opt/SUNWspro/lib/libw.so.1
28947: trying path=/usr/openwin/lib/libw.so.1
28947: trying path=/usr/dt/lib/libw.so.1
28947: trying path=/usr/local/lib/libw.so.1
28947: search path=/opt/SUNWspro/lib/rw7:/opt/SUNWspro/lib:/opt/SUNWspro/lib:/usr/ccs/lib:/usr/lib (RPATH from file merry1)
28947: trying path=/opt/SUNWspro/lib/rw7/libw.so.1
28947: trying path=/opt/SUNWspro/lib/libw.so.1
28947: trying path=/opt/SUNWspro/lib/libw.so.1
28947: trying path=/usr/ccs/lib/libw.so.1
28947: trying path=/usr/lib/libw.so.1
28947:
28947: find library=libc.so.1; searching
28947: search path=/opt/SUNWspro/lib:/usr/openwin/lib:/usr/dt/lib:/usr/local/lib (LD_LIBRARY_PATH)
28947: trying path=/opt/SUNWspro/lib/libc.so.1
28947: trying path=/usr/openwin/lib/libc.so.1
28947: trying path=/usr/dt/lib/libc.so.1
28947: trying path=/usr/local/lib/libc.so.1
28947: search path=/opt/SUNWspro/lib/rw7:/opt/SUNWspro/lib:/opt/SUNWspro/lib:/usr/ccs/lib:/usr/lib (RPATH from file merry1)
28947: trying path=/opt/SUNWspro/lib/rw7/libc.so.1
28947: trying path=/opt/SUNWspro/lib/libc.so.1
28947: trying path=/opt/SUNWspro/lib/libc.so.1
28947: trying path=/usr/ccs/lib/libc.so.1
28947: trying path=/usr/lib/libc.so.1
28947:
28947: find library=libdl.so.1; searching
28947: search path=/opt/SUNWspro/lib:/usr/openwin/lib:/usr/dt/lib:/usr/local/lib (LD_LIBRARY_PATH)
28947: trying path=/opt/SUNWspro/lib/libdl.so.1
28947: trying path=/usr/openwin/lib/libdl.so.1
28947: trying path=/usr/dt/lib/libdl.so.1
28947: trying path=/usr/local/lib/libdl.so.1
28947: search path=/usr/lib (default)
28947: trying path=/usr/lib/libdl.so.1
```

```

28947:
28947: calling .init (from sorted order): /usr/lib/libc.so.1
28947:
28947: calling .init (dynamically triggered): /usr/lib/libCrun.so.1
28947:
28947: calling .init (done): /usr/lib/libCrun.so.1
28947:
28947: calling .init (done): /usr/lib/libc.so.1
28947:
28947: transferring control: merry1
28947:
Merry Xmas, Happy 2002!
28947:
28947: calling .fini: /usr/lib/libCrun.so.1
28947:
28947: calling .fini: /usr/lib/libc.so.1

```

Error messages about missing libraries are usually due to the fact that a shared library is placed in a nonstandard location (not defined in `ld.so.conf`). In that case you should add the directory containing the supplementary library to the `LD_LIBRARY_PATH` environment variable. However, in that case you also should have specified the same path with the `-L` option on the `g++` command to verify that all externals are resolved by the libraries at link time.

As libraries are loaded, the system must occasionally perform certain preliminary steps. For example, in C++ programs some objects might have to be created before program startup, and destroyed before a program exists. This situation is handled by the dynamic linker by looking for `init()` (`fini()`) functions, that are created by the compiler and contain the necessary code needed to initialize (destroy and clean up) objects and other parts of the runtime environment. Calls to these initialization and finalization are clearly seen in the examples with the `LD_DEBUG` variable set to `libs`, shown previously.

Using shared libraries at CERN

Although both archive and shared libraries are offered for most packages, linking with the shared libraries is the default method on Linux with the `g++` compiler and linker/loader. However, the shared libraries for HEP applications (e.g., CLHEP, NAGC, Geant4, etc.) are not usually in any of the "standard system directories", so the path to these directories have to be communicated to the compiler/linker (`-L` option) at compile/link time and the dynamic linker (`LD_LIBRARY_PATH` environment variable) at run time. Presently, on Linux, there is the additional complication that the default version of `gcc` (2.91) is not adequate for many C++ applications and the preferred `gcc` 2.95 has to be specifically requested. For CERN applications there are usually well documented startup scripts (e.g., the "Setting up the user environment" of the Anaphe manual) which provide this functionality automatically.

The three examples that follow are set up to be run on the *standard Linux environment* at CERN (i.e., they have all been tested on `lxplus`). On Sun Solaris or other platforms similar but slightly different command names, setting of environment variables, library paths, etc. might be required, but the overall strategy should remain the same.

Using CLHEP

An example of using the CLHEP library follows. First we define a few environment variables to set the correct compiler version, and to point to the CLHEP hierarchy (in fact, as shown in the last example, one can get access more simply to the packages by initializing the Anaphe environment scripts).

```

# Define specific version of compiler libraries
$ export GCCVER=/usr/local/gcc-alt-2.95.2
# Include compiler in PATH
$ export PATH=$GCCVER/bin:$PATH
# Define which version of CHLEP we want to use
$ export CLHEPVER=1.7.0.0
# Define where the platform/compiler specific files reside
$ export CLHEPSP=/afs/cern.ch/sw/lhcx/specific/redhat61/gcc-2.95.2/CLHEP/$CLHEPVER
# Compile the program
$ g++ -I$CLHEPSP/include -o testRandom.o -c testRandom.cc
# Define the path of the shared libraries
$ export LD_LIBRARY_PATH=$GCCVER/lib:$CLHEPSP/lib:$LD_LIBRARY_PATH
# Link the program
$ g++ -L$CLHEPSP/lib -lCLHEP -o testRandom testRandom.o
# Execute the program
$ testRandom
-0.715106 -0.0593743 0.276891 1.25028 0.126944

```

The C++ program `testRandom.cc` itself is shown next. It is an example of a random sequence of values according to a gaussian distribution, using a generator based in the Mersenne Twister engine algorithm. This procedure is often used in Geant4. The relative paths of the include files on lines 1 and 2 are dereferenced with the help of the `-I` option of the `g++` command, which specifies the top of the directory containing these include files.

```
#include "CLHEP/Random/Randomize.h"
#include "CLHEP/config/iostream.h"

int main(int, char**)
{
    HepInt size=5;
    HepDouble vect[size];

    // Create an instance of the algorithm object to be used
    // as underlying engine for the random generator
    //
    MTWistEngine* aRandomEngine = new MTWistEngine();

    // Instruct the generator to set the engine and use it for
    // for the generator of pseudo-random numbers
    //
    HepRandom::setTheEngine(aRandomEngine);

    // Fill the array 'vect' of random values generated
    // according to a Gaussian distribution centered on zero.
    //
    RandGauss::shootArray(size,vect);

    // Print the values on standard output
    //
    for (unsigned int i=0; i<size; ++i )
        HepStd::cout << " " << vect[i];
    std::cout << std::endl << std::endl;

    // Delete the engine object
    //
    delete aRandomEngine;

    return 0;
}
```

Running with CLHEP and NAGLIB

We show in this section how to link procedures both from the C library NAGLIB and from the C++ library CLHEP. Consider the following C++ program `test.cc`:

```
/* Test program tst.cc */
#include <iostream>
#include "CLHEP/Vector/ThreeVector.h"

#include <nag.h>
#include <nag_stdlib.h>
#include <nagg05.h>

int main(int, char**)
{
    std::cout << "Clhep test output" << std::endl;
    Hep3Vector v(1.,2.,3.);
    std::cout << v << std::endl;

    Int i;
    Int seed = 0;
```

```
std::cout << "g05cac Example Program Results" << std::endl;
g05cbc(seed);
for (i=1; i<=5; i++)
std::cout << g05cac() << std::endl;
return 0;
}
```

To compile and link this with gcc 2.95 on lxplus, the location of the compiler and package include files has to be defined, e.g., in tcsh:

```
$ setenv G295 /usr/local/gcc-alt-2.95.2
$ setenv PATH $G295/bin:${PATH}
$ setenv NAG /afs/cern.ch/sw/lhcxx/specific/i386_redhat61/Nag_C/6.0
$ setenv CLDIR /afs/cern.ch/sw/lhcxx/specific/i386_redhat61/CLHEP/new
```

Then the compile and link step producing the default a.out is:

```
g++ -I$CLDIR/include -I$NAG/include -L$CLDIR/lib -LCLHEP -L$NAG -lnagc -lm -lpthread tst.cc
```

In order to run the application, the directories containing the CLHEP, NAGC and gcc 2.95 libraries have to be given in the LD_LIBRARY_PATH environment variable. The ldd command shows how the libraries are dereferenced.

```
setenv LD_LIBRARY_PATH $G295/lib:$CLDIR/lib:$NAG
ldd a.out
libCLHEP.so => /afs/cern.ch/sw/lhcxx/specific/i386_redhat61/CLHEP/new/lib/libCLHEP.so (0x2aac0000)
libnagc.so.6 => /afs/cern.ch/sw/lhcxx/specific/i386_redhat61/Nag_C/6.0/libnagc.so.6 (0x2ab1e000)
libpthread.so.0 => /lib/libpthread.so.0 (0x2aeae000)
libstdc++-libc6.1-2.so.3 => /usr/local/gcc-alt-2.95.2/lib/libstdc++-libc6.1-2.so.3 (0x2aec2000)
libm.so.6 => /lib/libm.so.6 (0x2af09000)
libc.so.6 => /lib/libc.so.6 (0x2af26000)
libstdc++-libc6.1-1.so.2 => /usr/lib/libstdc++-libc6.1-1.so.2 (0x2b01b000)
/lib/ld-linux.so.2 => /lib/ld-linux.so.2 (0x2aaab000)
```

Finally, running the program produces:

```
./a.out
Clhep test output
(1,2,3)
g05cac Example Program Results
0.795124
0.225717
0.37128
0.225035
0.878745
```

Predefined package environment variables

Software packages developed at CERN come with setup scripts that define the relevant paths and environment variables needed to use the libraries referenced by the application in question.

As an example, for Anaphe (<http://anaphe.web.cern.ch/anaphe/>), both C and Bourne startup scripts are available that users have to source to define all needed environment variables. For instance with the variable LHCXXTOP pointing to the top of the Anaphe hierarchy one could add the following in the user's .profile file (Bourne shell):

```
export PATH=/usr/local/gcc-alt-2.95.2/bin:$PATH
export LD_LIBRARY_PATH=/usr/local/gcc-alt-2.95.2/lib:$LD_LIBRARY_PATH
export LHCXXTOP=/afs/cern.ch/sw/lhcxx
# Source the startup script. If using versions other than 3.6.3
# just substitute it with the new version.
VER=3.6.3
. $LHCXXTOP/share/LHCXX/$VER/install/sharedstart.sh
```

The VER variable is the version of the software to be used (e.g., 3.6.3 November 2001).

Let us now take part of the Anaphe suite, and see how this works. We choose an example of the fitting and minimization package FML, a C++ library with a rich set of classes. The files from the stlfit source directory were copied locally (/afs/cern.ch/sw/lhcxx/share/FML/1.2.0.1/fml/examples/stlfit). Then, we used the following GNUmakefile (a slight variant of the file provided in the original source distribution):

```
include $(LHCXXTOP)/share/FML/1.2.0.1/common/fml-app.mk

# if you want to compile with minuit type: gmake "USE_MINUIT=1"
# by default nagc is used

all:    stlfit

stlfit: stlfit.cpp
        $(C++) $(C_FLAGS) stlfit.cpp -o stlfit -L$(LHCXX_REL_DIR)/lib $(LIBS_FML) $(LIBS_VECTOR) -lCLHEP $(DYNA_LIBS)

clean:
        rm -f stlfit
```

We prepare the executable by running gmake:

```
$ gmake "USE_RELEASE=1"

g++ -ansi -pedantic -Wall -Wno-long-long -Wpointer-arith
    -Wnested-externs -Woverloaded-virtual -Wbad-function-cast
    -Wconversion -fnonnull-objects -pthread -I
    /afs/cern.ch/sw/lhcxx/specific/redhat61/gcc-2.95.2/3.6.3/include
    stlfit.cpp -o stlfit
    -L/afs/cern.ch/sw/lhcxx/specific/redhat61/gcc-2.95.2/3.6.3/lib
    -lFML-nag -lGemini-nag -lnagc -lpthread -lVector -lAIDA_Annotation
    -lHepUtilities -lCLHEP -ldl
```

We see above that the correct list of libraries, include files, etc. is declared automatically thanks to using the provided GNUmakefile. Finally, we run the example, as follows:

```
$ stlfit

default constructor

reading in 1D-vector of size 10 (10, )
-----
ExtendedFitter setup

Number of points: 9

Startup configuration
```

Index	Name	Value	Lower	Upper	Step
0	amp	-1.598e-01	no	no	1.000e+00
1	mean	-6.056e-01	no	no	1.000e+00
2	sigma	-2.169e-01	no	no	1.000e+00

```

fitting successful!

ExtendedFitter results:
Estimator = 1.057e+01
```

Index	Name	Value	Error	Minos+	Minos-
0	amp	-1.567e-01	2.824e+03	no	no
1	mean	-6.198e-01	1.338e+04	no	no
2	sigma	-1.371e-01	7.669e+04	no	no

```

-----
```

ExtendedFitter setup

Number of points: 9

Startup configuration

Index	Name	Value	Lower	Upper	Step
0	amp	-2.016e-01	no	no	1.000e+00
1	slope	-8.835e-02	no	no	1.000e+00

fitting successful!

ExtendedFitter results:

Estimator = 7.842e-12

Index	Name	Value	Error	Minos+	Minos-
0	amp	2.220e+00	5.316e-01	no	no
1	slope	-1.250e+00	4.138e-01	no	no

Every non-trivial software development effort should set up a similar working environment, whereby sourcing a script will define the relevant paths and libraries to allow one to run programs in a transparent way.

Constructing shared libraries

In our various examples we have repeatedly used shared libraries. However we have never shown in detail how to create these. Although this is not too complex in principle for C and C++ source code, the procedure is somewhat involved and system-dependent.

First, when shared libraries and dynamically loadable modules are compiled, often special compiler options have to be used for creating position-independent code (`-fpic`, `-FPIC`, `-KPIC`, and so on). These options tell the compiler to generate code that accesses symbols through a collection of indirection registers and global offset tables. The primary benefit of PIC code is that the text segment of the resulting library (containing machine instructions) can quickly relocate to any memory address without code patches. This improves program startup time and is important for libraries that large numbers of running processes must share, although on the other hand PIC can generate a modest degradation in application performance of a few percent. However, many shared library examples do not need PIC code for their shared libraries and dynamically loadable modules. Hence, if performance is important for a library or dynamically loadable module, you can compile it as non-PIC code. The primary downside to compiling the module as non-PIC is that loading time increases because the dynamic linker must make a large number of code patches when binding symbols.

Another possible issue while constructing a shared library is the inclusion of static libraries (`.a` files) when building dynamically loadable modules. Indeed, when parts of an application are built as static libraries, the linking process might not do what one expects, since the relevant parts of the static library are simply copied into the resulting shared library object. Thus, when these modules load, all private library copies remain isolated from each other. This can result in grossly erratic program behavior, e.g., with changes to global variables not affecting other modules, functions seeming to operate on different sets of data, and so on. In such a case, static libraries must be converted to shared libraries. Then, when multiple dynamic modules link again to the library, the dynamic linker will guarantee that only one copy of the library loads into memory.

Conclusion

We have explained how to compile, link and load C++ programs (the issues are similar for other compiled languages) and pointed out the advantages of shared with respect to static libraries. In particular, with the help of examples, we have introduced the major options of the `g++` command to specify libraries against which a program has to be linked. We also underlined the importance of specifying the correct `LD_LIBRARY_PATH` environment variable at load time. Finally, we mentioned how software packages, such as Anaphe and Geant4, assist their users by providing them with setup scripts and documentation to allow them work in the correct software environment automatically.

More detailed explanations, as well as a discussion of more advanced features, such as setting up filters, or loading new libraries at run time, are available in the original publication (see URL at beginning of article), which we strongly encourage everybody to read, as it serves as a useful complement to this text. We would like to thank Gabriele Cosmo and Jakub Moscicki for their help while preparing the examples, and Andreas Pfeiffer for several suggestions that improved the accuracy of the text.



8.2 Tips & Techniques from Users to Users

Compiled by IT/User Services

Introduction

This new column is aimed at sharing "tips and techniques" within the user community. If you have some practical advice to share, then we want to hear from you! Just send an e-mail to the CNL editor (CNL.Editor@cern.ch) for publication both in the Newsletter and on one of our forthcoming web pages.

The following "Useful links" and "tips" have been provided by a CERN user. The "tips" have been checked, and a few comments (in *Italic* font) have been included by the service responsables.

Some Useful links

1. *Mail tools* on the Web: change password, set holiday message, etc. (<http://mailwww.cern.ch/>).
2. List of *Common Acronyms and Abbreviations* encountered in the CERN environment (http://cern.ch/library/information_sources/accr.html).
3. *Administrative Procedures Manual* (<http://as.cern.ch/AdminMan/chapeng.html>).
4. *Helpdesk FAQ* (<http://consult.cern.ch/qas>).
5. *CERN Conversion Service*: Conversion between many different files formats, e.g., Microsoft Word to Postscript or Portable Document Format (PDF), $\text{\LaTeX}$ to HTML, gif and jpeg to PDF, etc. (<http://documents.cern.ch/Convert>).

Tips for Microsoft Word users

"Fast Save" option

Users should disable the "Fast Save" option to avoid unnecessarily large files.

For information, the "Fast Save" option saves *all* changes, so if a document has been edited a lot, even if it is one page, it can go up to 400 KB or more! In particular, it will bounce from most listboxes and is ridiculously large to send by email.

To deselect this option do the following:

"Tools" menu,
"Options",
"Save" tab
deselect "Allow fast saves"

Please note that the default installation of MS Word leaves the "Fast Save" option disabled and the NICE team recommends to keep it disabled. Furthermore, "Fast Save" cannot be performed when working on a document over the network.

Naming of documents

When sending e-mails you should **always** use the correct extension names, e.g., word.doc, excel.xls, powerpoint.ppt, etc. This will allow users to click straight away on documents to open them, without the need to re-save and rename them.

Specifying the correct extension is certainly a good practice, in particular it allows you to immediately see what the file type in attachment is. Note, However, that in general, Windows can recognize the file type and use the right tool to open the file, even without extension.

Page setup

You should verify that you have set the default page setup to A4. The procedure is the following:

"File" menu
"Page setup"
"Paper size" tab
select A4
click on default
answer YES

This will set A4 as the default page setup.

8.3 Questions and Answers from the Computing Helpdesk

Collected by the IT/User Services Team

Nicole Cr  mel (Editor), IT/User Services

Abstract

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

N.B. The number in parentheses refers to its relative num-

bering in the User Services "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 Services)

Question [542] – Extract web page contents Word (IE Kiosk mode)

I would like to implement a few web pages that I downloaded from the Internet, into a Word document.

Now, the only 2 ways I found out how to do this, is either by choosing select - copy (which chooses the whole document), or to take a screen shot, which takes exactly the page I want, but this leaves some not very nice header lines like the URL and the tool-bars on top of the image.

Is there an easy way to just select the page itself, without web address and toolbar, i.e. exactly the selected picture/text size and really nothing else?

Answer

If you use Internet Explorer (IE, currently v.5) there is a smart feature which does exactly what you want.

Internet Explorer has a full-screen mode that you access by pressing F11. However, the screen still shows boxes to minimise, resize, or close the window, so this is not exactly what you want either.

The KIOSK mode, on the other hand, hides all buttons. To start IE in Kiosk mode, from the Run prompt (Start - Run) start IE with the following command:

`ieexplore -k [site]`

where [site] is:

- blank: this uses your default start page;
- a URL, e.g., `http://consult/xfind`;
- a local Web page, e.g., `C:\folder\page.htm`;
- a remote Web page, e.g., `\\server\share\folder\page.htm`.

Within Kiosk mode, you can use these main commands (print them out beforehand, as afterwards you will be without any toolbar help!):

Ctrl+A: Select all (editing)
 Ctrl+B: Organise favorites
 Ctrl+C: Copy (editing)
 Ctrl+F: Find (on current page)
 Ctrl+H: View History folder
 Ctrl+L: Open Location dialog box (same as Ctrl+O)
 Ctrl+N: New window (this one opens in non-Kiosk mode)
 Ctrl+O: Open Location dialog box (same as Ctrl+L)
 Ctrl+P: Print
 Ctrl+R: Refresh
 Ctrl+S: Save
 Ctrl+V: Paste (editing)
 Ctrl+W: Close (same as Alt+F4)
 Ctrl+X: Cut (editing)
 Alt+F4: Close (same as Ctrl+W)
 Alt+Left Arrow: Back
 Alt+Right Arrow: Forward
 Esc: Stop
 F5: Refresh

The rest is a piece of cake: "printscreen" or "ALT+printscreen", to exit Kiosk mode (see above), paste the beauty into Word, and *bingo!*

Question [548] – Can I schedule a document to be printed overnight from a PC?

I have to print a huge document, but I don't want to monopolise our printer for something that might print for an hour or more, especially as it isn't urgent, nor confidential.

Is there a way to tell my PC to schedule it for overnight printing?

Answer

You can configure a given printer from your PC to only select your jobs between certain hours:

```
My Computer
Printers
Right-click on the printer you want
select "properties"
Click the "Scheduling" tab,
Click at top in the "Available" section
choose "From"
enter 'From' and 'To' time
(e.g., 18:00 - 08:30)
Click OK to save your changes
```

Jobs submitted by you to this printer from this PC will now only be printed between the hours specified, but be aware that this of course affects all your documents. However, your documents that are spooled from another PC (but the same account) will be printed normally.

Don't forget to change the settings back to "always" in the "available" field after your document has printed.

Question [554] – Hardware (MAC) address of a Windows 2000 PC

How can I find out the hardware address of my PC running under Windows 2000?

Answer

YOU can get the Media Access Control (MAC) address of your W2k PC as follows:

- right-click "My Computer"
- click "Manage" and then:
 - (+) System Tools
 - (+) System Information
 - (+) Components
 - (+) Network > Adapter > MAC Address

[the (+) sign means expand the tree]

Question [564] – Profile conflict NT / W2000

By mistake, I have logged into an NT4 PC and afterwards into one with OS W2000. Now this computer warns me about possible profile conflicts.

Answer

Although in many cases not dangerous, the message is at least bothering and just tells the user there *might* be conflicts.

To get rid of it, please literally (really!) follow the procedure as described, including a scaring statement saying "just power off your PC". This message has to do with the fact that with a normal, proper rundown the file `ntuser.dat` gets

endlessly re-written back to the server, which in the present case is obviously not desirable.

And bear in mind: *Never* use the same NICE account on NT and W2000 in parallel, but get 2 different accounts, one for each OS. This is recommended in the *Win Services Documentation* at the URL <http://winservices.cern.ch>, and in the related QA at the URL <http://consult.cern.ch/qa/0576>.

Question [582] – Access rights to a single NICE file (not folder)

I use the trustee manager (<http://winservices>) to assign rights to some of my folders, but haven't found a way to assign rights to a single file within a folder.

Answer

It was officially decided at CERN that the NICE trustee manager only lets users modify rights for folders, but not files.

The only workaround is therefore to create a new folder, put the file into it and then give rights to this folder.

Question [590] – NICE group manager structure

I am registered as the owner of a divisional sub-group (say IT-XX) but don't seem to have any rights on it, e.g., adding another user via the NICE group manager accessible from <http://winservices>.

Answer

Being registered as the owner of a group doesn't necessarily mean you have full rights.

The structure of the NICE group manager is:

```
Group manager
--contains--> Classes
--contain--> Groups
--contain--> Members
```

So, in order to have full rights of IT-XX you should have been registered on the next higher level, in this case in the class IT, which you aren't. More, you don't even want to be, as this would empower you to have rights over other groups in there which you don't need.

A general solution is therefore (same example taken again):

- Create a class IT-XX Admins,
- and create an associated group IT-XX Users.

Question [592] – Install applications from dfs

I am trying to install an application I found on dfs (<http://cern.ch/dfs>) for my Windows 2000 PC, but need further advice.

Answer

Applications you only find on *dfs* are not supported by the NICE 2000 team, they are just there *as-is*. Applications you need support for must only be installed from the control panel using "add/remove programs - add new programs"

Question [603] – Audit of hardware and software features of a W2000 PC

Full audit on my W2000 hardware and software functions, features, limitations, etc.

Answer

W2000 allows you to run a full or partial audit of practically all hardware and software "bells and whistles" your PC is capable of.

The command `msinfo32 /?` gives you a brief overview over the command syntax. Thus, for instance, you could run a total audit of your PC configuration using the command:

```
msinfo32 /report C:\report.txt
```

This outputs its report onto the file `report.txt` of your C\ drive.

Note that during this procedure (which takes a few minutes) the above file can not be viewed. Also, as all hard- and software features are queried on a full report, your local disk(s) might make some strange vibrating noise for a few seconds.

Question [636] – Office 2000 detect and repair feature

I keep having constant problems with a W2000 Office application.

Answer

Although Office 2000 applications are very stable, you may (apart from the "usual" recipes) also try to repair a particular W2000 Office installation.

End-users can remove, change Office components only if they have sufficient privileges to perform these tasks. In practice, only the member of the local Administrators group can do this. However, any user can repair any office 2000 installation.

To repair Office, just select "Detect and Repair" from the "help" menu within any Office application.

Question [672] – Euro symbol with Office

Is the Euro symbol available on Office 97 and 2000?

Answer

Yes. There are, in fact, different possibilities to use the Euro symbol inside both Office 97 and Office 2000 applications:

Word and Powerpoint have the Euro symbol in their insert - symbol collection. Excel, however, is a bit more complicated, as it doesn't have this function. To type it inside Excel, use the key combination "ALTGR+0128 (where ALTGR is the right/hand ALT key) on a US keyboard. This key sequence is also available on all other Office applications.

Users who have a French keyboard must use the combination "ALTGR+E" instead.

To apply the Euro currency format to values in your Excel worksheet, first select the values you want to format. On the Format menu, click Cells, and then click the Number tab. In the Category list, click Currency, and then click the Euro currency format you want to use in the Symbol list.

There is, of course, also the possibility to copy/paste the Euro symbol. On NICE 95/NT the procedure is:

More Applications

- Accessories
- CharMap

On NICE 2000/NT the procedure is:

Programs

- Accessories
- System Tools
- Character Map

UNIX Support (General)

Question [529] – ls: No such file or directory (with AFS)

The UNIX command `ls`, for a directory that does exist (I can see it if I login on another node), gives the error message:

```
ls: /afs/cern.ch/atlas/project/graphics/DB/Linux:
No such file or directory
```

What is the problem?

Answer

This behaviour usually indicates an AFS corrupted entry for this directory, while it is still visible on other nodes. Normally this error will timeout within the next 24 hours, otherwise the machine node (with the problem) will have to be rebooted by the service managers.

Meanwhile users should logout and login on another node.

Question [544] – Scratch area on ATLAS

- Variant 1
I used to work on the machines `atlas.cern.ch`, where I put my data files in the directory `/scratch/zp/username/...`. As there is no access anymore to this `/scratch` folder on the LXPLUS machines what shall I do with my data files?
- Variant 2
I am a user of ATLAS and would like to get some scratch area for my analysis jobs. Whom shall I contact?

Answer

The ATLAS interactive nodes are being reinstalled with a more standard configuration and hostname, in preparation for a merge into the LXPLUS cluster. The major difference users will see is the removal of NFS accessible file-systems, which was used for the scratch area. This was agreed between the IT and ATLAS computing coordinators, for issues of system management and maintenance.

The functional alternative suggested is AFS scratch areas assigned under the user home directories as `scratch0`, `scratch1`, etc. These directories are personal to the account so must be managed by the owner and, being scratch, are not backed up. Users can copy existing files to these new directories by logging in to the old ATLAS nodes.

Currently we invite all users to forward directly such requests to AFS support (Afs.Support@cern.ch) to get the AFS area setup, specifying how much space (in Megabytes) they think they need.

Please give feedback on the usability of this new solution via the ATLAS mailing list, at atlas-support@listbox.cern.ch.

UNIX Support (Batch)

Question [547] – batch jobs - No AFS token

It seems that when too many batch jobs try to start at the same time, it can happen that some jobs die with the error message

```
No AFS token
@(#)CERN job starter 2.08 6/6/2000
```

Answer

The latest version of the `job_starter` should try to avoid this behaviour (sleeping for a few minutes and try again later), but it does not prevent it at 100%. A general advice is, when job production managers start big production, to insert “sleep 60” between two `bsub` commands in the script that submit the series of jobs.

Scientific Software

(Nicole Crémel, IT/User Services)

Question [522] – FORTRAN 90 on LXPLUS

What is the situation with FORTRAN 90 compilers on the LXPLUS Unix central service?

Answer

Here is the “official” statement, as of March 2001, made by the Group Leader of IT/API Group (Applications for Physics & Infrastructure).

As you may know, CERN is decreasing the FORTRAN support significantly and, e.g., CERNLIB will have its last release in 2002 and be supported in a limited way for two years after that. Several years ago it was decided that we would not move to FORTRAN 90 neither for CERNLIB nor for compiler support in general. The experiments have agreed to this strategy and are actively moving towards C++. For limited and specialized use we do provide the PGF90 compiler as you mentioned. The access and use is described at:

<http://wwwinfo.cern.ch/pdp/is/eric/pgf.txt>

We have no resources to extend this service in any way—in particular not by providing additional compilers from other suppliers.

This policy towards C++ has been agreed by the FOCUS committee: (<http://cern.web.cern.ch/CERN/meetings/FOCUS/Welcome.html> which is also the place to formulate requests that are against this official statement (Mailing list:

cern-focus@listbox.cern.ch).

I regret not to be able to give you a more positive answer and remain.

With best regards – Jürgen Knobloch

Question [546] – Linking f90 code with CERNLIB (e.g., graflib) on Solaris

We have to port an application using the HIGZ graphics package from CERNLIB to a Solaris platform. This code was originally developed on Digital UNIX in Fortran 90. Trying to link this application with f90 and the `cernlib` command we get many undefined symbols at link time:

```
'f90 -o zzz xxx.f90 'cernlib graflib'
Undefined                          first referenced
 symbol                            in file
__e_wsle                           /cern/pro/lib/libgrafX11.a(iopks.o)
__s_stop                           /cern/pro/lib/libpacklib.a(zend.o)
__do_f_out_nv                       /cern/pro/lib/libpacklib.a(mzdiv.o)
__s_wsle_nv                        /cern/pro/lib/libgrafX11.a(iopks.o)
etc...
d: fatal: Symbol referencing errors.
      No output written to zzz
```

Answer

Following is the answer provided by the CERN Program Librarian. However please note that Fortran 90 is not supported at CERN, as explained above and at the URL <http://consult.cern.ch/qa/522>.

This is a bit tricky. The missing externals are from the f77 runtime library which is not included by default by the f90 and f95 Solaris WS6 compiler invocations. However, there is an option to simplify inclusion of the f77 compatibility library `-xlang=f77`, which I found to work for simple applications. So try:

```
f90 -xlang=f77 -o zzz xxx.f90 'cernlib graflib'
```

When the option `-xlang=f77` is not recognised by the compiler, the solution can be to find and add the specific f77 runtime library for the system. E.g., on Solaris, `/opt/SUNWsp/ro/lib/libf77compat.so`, try:

```
f90 -o zzz xxx.f90 'cernlib graflib' \
/opt/SUNWsp/ro/lib/libf77compat.so
```

Question [552] – commands xgrab and xgrabsc on LX-PLUS (Linux)

I am a great friend of the commands `xgrab` and `xgrabsc` on UNIX. It seems to me that this doesn't exist on LXPLUS? Is there something equivalent?

Answer

Unfortunately those commands are not available on LINUX, as you can see in the ASIS page:

<http://asis.web.cern.ch/asis/products/X11/xgrabsc.html>

The ASIS team is inviting users to use instead either `import` (a command from *X11/ImageMagick-4.1.6*, see the URL <http://asis.web.cern.ch/asis/products/X11/ImageMagick.html>), or `xv`, and `xwd` or `xpr` that work fine on Linux.

Question [556] – Visual C++ on a PC/Windows

I have problems/question with a C++ application developed with Visual C++ on a PC/Windows.

Answer

Visual Studio is supported on NICE only in the sense that they will ensure that proper installation can be made. There is no programming or debugging support provided by the NICE team.

IT policy is to support development of scientific software only on Unix, and more precisely only Linux and Solaris, so there is **no support** for software development (Visual Fortran, C++, etc.) on Windows/NICE. Only the LHCb experiment decided to build their development on Windows, but they are supporting themselves.

In practice this means that questions in this area submitted to IT services (including the Helpdesk) will get a reply showing this standard answer. When choosing to develop on a Windows platform users must understand that they are mainly on their own for software development issues.

Please note: it should also be understood that only the programming environment and infrastructure (compilers, loaders, libraries, etc.) for software development will be supported on Unix (Linux and Solaris), but we cannot afford to give general computing usage support for any of the languages. For this, users should refer to specialized books.

Mail Support**Question [525] – Netscape mail password (remember or not)**

I have set my Netscape mail client to always remember the mail password when I access my mail from an account other than my own. How can I revert, i.e., tell Netscape mail to always prompt me for a password?

Answer

From the Netscape session select:

```
Edit - Preferences...
Mail & Newsgroup
Mail Servers
select your mail server
click on button "Edit..."
untick "Remember Password"
```

Question [540] – Multiple Netscape address-books and bookmarks backup on NICE

I have multiple Netscape address books on my PC and lost all of them by some bad manipulation. Who can restore them?

Answer

Within Netscape (currently v.4.75), it is possible to have multiple address-books (and also bookmark files). However, this is not the CERN standard configuration, and only the default one is mirrored to the user's NICE home server, whereas all others are just stored locally.

Therefore, only this default version,

```
'userid'\w95\nscape4\pab.na2'
```

for the address-book and

```
'userid'\w95\nscape4\bookmark.htm'
```

for the bookmark file, can be restored from a backup copy in case of file loss or corruption.

For this reason, users are strongly encouraged to stick to the CERN standard Netscape setup.

Question [579] – Difference between Outlook and Outlook Express

What's the difference between Outlook and Outlook Express?

Answer

Outlook Express is an e-mail client and newsreader which comes for free with Internet Explorer. That's just about it.

Outlook (current version 2000) is part of the M/S Office package and is therefore not free.

Not only has it the facilities of Outlook Express, but it is also a good tool for some message networking, and contains modules for group work and a scheduler.

So, Outlook Express is a basic tool for a single PC, whereas Outlook 2000 is more sophisticated and desirable for LANs, networking functions, the like.

Text Processing

Question [539] – How can I edit an Spanish document in L^AT_EX?

Is it possible to create at CERN L^AT_EX documents in Spanish? In particular I am interested in accessing L^AT_EX packages to enable proper hyphenation and spelling dictionaries.

Answer

At CERN, there exists a command in the ASIS repository which sets up T_EX for basic Spanish typography, namely `eslatex`. It is just an appropriate call to `tex`.

Supposing you have Spanish hyphenation patterns compiled into your L^AT_EX format, you can also explicitly insert Spanish-related instructions in your document. For instance:

```
%Specify Spanish as main language
\usepackage[activeacute,spanish]{babel}
% Allows accented characters as input
% This is the default on Unix
\usepackage[latin1]{inputenc}
```

More about Spanish typesetting with L^AT_EX can be found at the Website <http://apolo.us.es/CervanTeX/>

Miscellaneous

Question [545] – 192.91... TCP/IP addresses at CERN

I work on a computer at CERN whose IP address is 192.91.xx.xx. Do I have standard network support for it, in case of problems?

Answer

TCP/IP addresses starting with 192.91... are being used at CERN, but in all cases point to a privately set up and owned network (private Cisco's and the like). The owner of this computer (see: <http://cern.ch/register>) has the responsibility and should be contacted.

Question [550] – Problem or question on CDR (Central Data Recording)

I have a problem (question) on the CERN Central Data Recording system (CDR)

Answer

Look at the CDR FAQs database, as provided by the CDR Support team, at the URL: <http://cern.ch/cdr/faq>

Question [562] – ftp to CMS (SUN / Solaris machines)

When I try to establish an ftp session from anywhere to CMS (e.g., from a NICE PC) I get the following message:

421 Service not available, remote server not reachable

Answer

For security reasons system responsible had to close the ftp daemons on SUN machines. Users should use `ftp` to LX-PLUS instead. If the data is local to CMS, use `scp` to copy it to LXPLUS.

If you have an account on the CMS cluster the only ftp server you can now reach from externally is `cmsftp01.cern.ch`

About the author(s): *Nicole Crémel is working in the User Services 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, offerings of the Computing Bookshop, etc.

9.1 News from the User Services Bookshop

Roger Woolnough and Jutta Megies, IT/User Services

This year the Frankfurt Bookfair was again punctuated with many takeovers and consolidations of companies. *Elsevier* has taken control of *Academic Press/Harcourt* and *Butterworth Heinmann/Morgan Kaufmann*. This makes for a very large Group and also will result in a virtual monopoly of Journals. The other big takeover is for *Wiley* who have acquired *Hungry Minds* (ex *IDG*). This should have a positive effect for our Bookstore as we should be able to negotiate a better rebate than we have at present with *Wiley*.

As far as new acquisitions are concerned the most recent revolve around dictionaries. We will be selling an Oxford “pop up” dictionary that can be loaded onto a PC and will trans-

late into English/French/German & Spanish as required. Just before Christmas the new English/French dictionary will be available, more than 900,000 words, phrases and translations (*Oxford Hachette* edition).

Of interest and arriving end January 2002 is the new *Cambridge University Press* book, *Numerical Recipes in C++*. We will send a news item when this book is available.

We have new books on PHP, XML, ColdFusion and a new Java Cookbook. On the general side we are pleased to announce the *Feynman Lectures on Computation* which is now available.

9.2 User Services Book Catalogue

Roger Woolnough and Jutta Megies, IT/User Services

Situated in Building 513 1-022 (tel. 74050) is the Computing Bookshop (<http://consult.cern.ch/xwho/people/67659>) provided by the *IT / User Services* group, where CERN users can find computer books and CDs at discount prices. The service is open weekdays from 8.30 to 12.30 or contactable by e-mail to Bookshop@cern.ch.

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

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

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

10. Just For Fun ...

Some jokes to end this CNL, and make more fun... for those who did not read them already, somewhere else!

10.1 Théorème de l'Euro (en français, in French)

Abstract

Il me semble opportun de mettre une note d'humour dans notre CNL, et de reproduire ce théorème d'actualité, que nous appellerons "le théorème de l'Euro". Cela peut intéresser tous les utilisateurs du CERN habitant la France voisine...

Le théorème est le suivant: "Avec l'Euro, c'est un peu plus cher que quand c'est moins cher, mais ça reste moins cher que quand c'est plus cher..."

Comme tout bon théorème doit pouvoir se démontrer, voici donc la démonstration de celui-ci, de façon strictement empirique, par une petite histoire qui pourrait nous arriver prochainement.

Le 1er janvier 2002, sans un seul Euro en poche mais le porte-monnaie plein de FF, je vais chez ma boulangère préférée lui présenter mes meilleurs vœux et acheter une baguette bien croustillante. Le dialogue est le suivant (avec quelques commentaires en italique):

Moi : A quel prix est la baguette maintenant ?

La boulangère : 4,30 FF comme d'habitude...

Moi : Oui, mais à partir d'aujourd'hui c'est en Euro.

La boulangère : Ah ! J'oubliais...

La boulangère sort son "EuroCalcullette" et fait le calcul.

Ca fait donc 0,66 Euro.

Moi : 0,66 Euro... mais au fait 0,66 Euro ça fait combien de FF ?

Je sors à mon tour mon EuroCalcullette...

Mais ça fait 4,33 FF ! La baguette a donc augmenté de 3 centimes ?

La boulangère : Non, c'est toujours le même prix...

Moi : Je comprend, c'est une "hausse mécanique" dis-je très "technocratiquement"

Si je vous donne une pièce de 5 FF, vous me rendez la monnaie en Euros ?

La boulangère : Tout à fait.

Ma boulangère est très "mode" et ne dit plus "oui".

Alors 5 FF, ça fait ...

Elle resort son EuroCalcullette...

0,76 Euro moins 0,66 Euro la baguette, je vous rends donc 10 centimes d'Euro.

Certains disent 10 cents, mais on peut confondre avec les US, alors comme c'est déjà pas si simple...

Moi : Ma première pièce en Euro ! et en FF c'est combien ? 0,10 Euro.

Je reprend mon EuroCalcullette...

0,66 FF,

tiens c'est comme le prix de la baguette en Euro

mais attendez !... 5 FF moins 0,66 FF, ça met le prix de la baguette à 4,34 FF ?

La boulangère : Non, non, c'est toujours le même prix.

Moi : Tout compte fait, je me demande si je ne devrais pas prendre 2 baguettes. Calculons: $4,30 \times 2 = 8,60$ FF, soit

après avoir repris mon EuroCalcullette...

1,31 Euro. 2 baguettes coûtent donc moins cher que 2 fois une baguette ? Et si je paie avec une pièce de 10 FF (= 1,52 Euro), moins 1,31 Euro, les 2 baguettes = 0,21 Euro, donc 1,38 FF, ce qui met le prix de la baguette à: $(10 - 1,38)/2 = 8,62$ FF soit 4,31 FF !

La conclusion est donc la suivante: la baguette qui coûte 4,30 FF coûte désormais

- 4,34 FF si vous l'achetez à l'unité
- 4,31 FF si vous les achetez par 2...

mais c'est quand même toujours le même prix. C'est donc : *un peu plus cher que quand c'est moins cher, mais ça reste moins cher que quand c'est plus cher !*

Notre théorème est ainsi démontré... En attendant vos prochaines expériences nous vous laissons à vos EuroCalcullettes.

