



# DESY Computing Newsletter No. 1 June 1991

A Publication of the DESY User Support Group

## Contents

• Editorial . . . . .	1
• The DESY User Consulting Office . . . . .	2
• New Central Computer installed . . . . .	4
• Why you should NOT use the Siemens Compiler . . . . .	6
• Remote Login . . . . .	7
• Usage of CERNLIB at DESY . . . . .	10
• FORTRAN Tips . . . . .	13
• Fortran 90 - the current state . . . . .	14
• The Database System DB2 at DESY . . . . .	15
• Questions and Answers from the UCO . . . . .	18

---



## ● Editorial

### Yet another Newsletter ?!

You have just obtained the first issue of the newly created quarterly *DESY Computing Newsletter*. The DCN (for those who like acronyms) is intended for all users and suppliers of data processing at DESY. We hope to offer

- General News
- Information articles
- Information on new offerings
- Miscellaneous notes
- Items from the UCO work

The DCN is produced by the User Support Group of -R-, with three persons from that group serving as editors. The major part of this issue was written by the editors themselves, but we hope that in further issues more articles from other authors will appear in the DCN. We would be happy to assist anybody who is willing to contribute to the DCN with advice and editorial assistance.

This first issue concentrates mostly on central computing because the editors are more familiar with this area than with the workstation or PC world. In future issues however, we will try to extend the scope to cover the full range of applications and environments.

We are starting with a very simple scheme for the distribution of the DCN: *We do not distribute it at all !* We supply copies of the DCN at the DESY computing center and the DESY library for self-service. We do not plan to mail personal copies to people on the DESY site.

People from outside who have no chance of getting a copy this way should contact

Mrs. Fock, ext. 3021

to obtain a copy. The e-mail address of Mrs. Fock is R01FOC at DHHDESY3.

We would be very happy to receive comments from our readership, telling us what is missing in the DCN and in which areas we should improve. We would be even happier if someone is willing to join us - be it regularly or temporarily. Please contact any of us, don't hesitate.

*The Editors*

#### Editors of the DCN are:

Katherine Wipf	R01WIP	ext. 3222
Michael Behrens	R01BEH	ext. 2556
Jan Hendrik Peters	R01NET	ext. 2583



## • The DESY User Consulting Office

*by Katherine Wipf*

On May 28<sup>th</sup>, 1990, the Desy User Consulting Office (UCO) opened its doors for the first time. The UCO has been offering its services for about a year now, and in this time it has seen several improvements:

- June, 1990: The User Support Group (USG) was formed to provide support for the users reaching beyond the services of the UCO.
- November, 1990: Walls were built around the UCO desk, creating a proper office for user consulting.
- At about the same time: The office started stocking the most commonly used CERN literature.
- April 25<sup>th</sup>, 1991: a PC was installed in the consulting office for transferring files from the IBM to diskette and vice versa. This service is now available during the UCO opening hours.

## Who We are and Where and When to find Us

The regular user consultants are Katherine Wipf, Michael Behrens, and Jan Hendrik Peters. Although each of us has an area of special expertise, you can come to us with questions outside these fields as well. Occasionally other colleagues from the computer center will substitute for us to ensure that the office is always open during the regular hours.

The UCO is located in room 129b of building 2B. This room is directly in the computer center (Rechenzentrum), in the terminal pool near the L1 printer. The opening hours are as follows:

Hours	Consultant	Special Expertise
Monday 14 - 16:30	Katherine Wipf	All areas + VAX/VMS
Tuesday 14 - 16:30	Katherine Wipf	All areas + VAX/VMS
Wednesday 14 - 16:30	Jan Hendrik Peters	All areas + E-Mail
Thursday 14 - 16:30	Katherine Wipf	All areas + VAX/VMS
Friday 9:30 - 12:00	Michael Behrens	All areas + Fortran, C, Pascal

During these hours you can telephone the office by dialing the extension 2510 (outside DESY you will need the full number: (040) 8998 2510). You can also contact the UCO via electronic mail by sending your message to:

R01UCO@DHHDESY3.BITNET

## What We Offer

### *Problem Solving*

The User Consulting Office should be the first place you turn to when you need help or assistance in using a computer at DESY.

We deal with problems ranging from "Why doesn't my job/program run?" and "Which software package should I use for inverting my matrix?" to "How can I transfer my VAX files to the IBM?" and "How do I login to DESY from other computers?" and many more. In most cases the consultant on duty can solve the problem immediately. If your problem requires a specialist, the consultant can tell you who to talk to or can call the expert for you. This should relieve you of the tedious job of locating a specialist for the particular problem you are having.

### *Documentation*

In addition to helping with computer problems, the UCO provides a number of manuals for users to pick up as the need arises. These include:



- IBM reference manuals for the programming languages VS FORTRAN, Pascal, and C .
- CERN documentation for the most commonly used CERN packages such as CMZ, PAW, MINUIT, HBOOK, HPLOT, CERNLIB, etc.
- Short reference summaries for various IBM products such as DB2, QMF, AS, Entry Assist, etc.
- "How to get Started on a VAX" by Karl Gather.

The UCO also provides a central ordering service for CERN manuals (by agreement with CERN), so if you need any personal copies of CERN documentation which the UCO does not have, please contact us.

The User Support Group itself writes documentation related to the special DESY environment. So far two papers have been published which are available in the computer center self service area and in the UCO:

- *J.H. Peters*, Printing at DESY, DESY-USG/90/01
- *M. Behrens*, IBM/MVS Extended Architecture, DESY-USG/90/02

Other papers will follow soon. Already in preparation are *Logging on at DESY*, *A Short Guide to PAW at DESY*, and *JCL for Beginners*, others have been started.

## *File Transfer Services*

The UCO's newest feature is a file transfer service between the IBM mainframe and  $3\frac{1}{2}$  or  $5\frac{1}{4}$  inch diskettes (IBM compatible). The UCO PC can read both single- and double-density diskettes. Please bring your own diskettes for transfers from the IBM mainframe to the PC.

## **Our Plans for the Future**

### *A Central Data-base for all DESY Computer Users*

The USG plans to build and maintain a central data-base which should contain everyone who has a userid on a DESY computer. For each user the data-base

will contain information like: full name, telephone number, home institute (for external users), and a list of all his DESY userids with the corresponding computers. Once this data-base has been created, the UCO will be able to supply inquirers with information on how to contact a DESY user, regardless of which machine he usually works on.

### *An Extended File Transfer Service*

We have ordered a TCP/IP card for the UCO PC and plan to install it as soon as it arrives. This will make it possible to transfer files to and from the PC diskette using FTP. As a result, the file transfer service can be extended to any machine with a TCP/IP implementation (i.e. VAXes, Workstations, PCs, etc.).

### *We Welcome Your Suggestions*

If there is a special service which you would like the User Consulting Office to offer, please come and talk to us. We cannot promise to offer all services or support which you wish for, but we are willing to consider any serious suggestions.



## • New Central Computer installed

*by Paul Kuhlmann*

On May 25<sup>th</sup>/26<sup>th</sup> 1991 - almost the same time as this first issue of the DCN - a new central computer has been installed at the DESY computing centre. The former machine, an IBM 3090-30S computer, has been upgraded to a model ES 9000-720. The old machine had 3 processors, the new one has 6 processors, each with the same processing power as the former ones. So the processing power is doubled with this upgrade. Central memory, expanded memory and the number of I/O channels are increased accordingly. The figure on the next page shows the new configuration.

This transition is accompanied by a whole series of changes, so that we can see this as a new phase for the computing centre:

- The new operating system MVS/ESA replaces the former MVS/XA.  
In addition to the 31-bit address space (ca. 2 GB) that was present with the XA version of the operating system, this new version allows multiple *data spaces* of the same size. This new feature will be used primarily by the operating system itself, opening a path to more effective I/O operations.
- The newly announced serial (fiber optic) channels can be used with the new machine. In the future this will drastically change the way I/O devices are connected to the central computer and will lead towards a more flexible configuration with increased data rates.
- As the old 3350 disk drives have now been sorted out, we are in the (rare) situation of running one single type of disk drive, the 3380. Unfortunately this will change again in the near future with the forthcoming installation of a new type of disk (named 3390) with increased performance but also with a different track length.
- The transition from magnetic tapes to cassettes is almost completed. Usage of magnetic tape

drives is now mainly for reading the existing archive tapes. Around 60 000 archive tapes exist today.

- The mass storage system(MSS) is not supported by the new machine. It was taken out of the system in May after 10 years of operation. Its function is now performed by using disks and the automatic cartridge silo (ACS).
- The new machine no longer needs special power supplies producing 440 Hz, so the old frequency transformers are no longer needed.
- The double floor in the computing centre had to be replaced partially before the installation of the new machine. It was installed in the early 70's, and its conductivity was no longer satisfactory for the new generation of machines.

## What does this all mean for a user of the computing centre?


We already made the transition to the new operating system a few weeks before the upgrade of the hardware, and we have had no negative effects so far. So we can forecast that our users will see no difference at all, except faster response times and increased 'kontingent' due to the increase in computing power.

The new possibilities of multiple address- and data spaces will for the moment only be available for system services, as their utilization requires privileged assembly routines.

Producing the new operating system, adapting it to our local needs and testing it thoroughly is a non-trivial task and it has kept the system programmers of -R1- busy for the last months. To do all this smoothly and without major interruptions of service is even less trivial and we are very happy with the results.



DESY-Rechenzentrum Stand Mai 91





## • Why you should NOT use the Siemens Compiler

by Michael Behrens

### The Status of the Siemens Compiler

The Siemens compiler for FORTRAN 77 was the first "standard" compiler for FORTRAN 77 at DESY as well as at many other IBM installations. The reasons for this choice was the poor quality of the first versions of the IBM compiler for FORTRAN 77. The IBM compiler has improved with time and now offers new functions, e.g. support for the vector processor, for large programs, interactive debug possibilities. As a result DESY decided to move to the IBM compiler (called VS FORTRAN) and to "freeze" the Siemens compiler. Most other installations did the same.

In the meantime, usage of the Siemens compiler has decreased considerably, especially for the new HERA experiments, but there is still a reasonable user community using it ( about 20% of all compilations in 1990). The situation is that the Siemens compiler

- has no XA - capabilities
- has no vector support
- will get no maintenance
- will get no new versions of libraries, including CERNLIB's ...

This means that Siemens users are cut off from new developments, that they will get no support for correcting compiler bugs, and that some day the compiler might not even work in an ever changing environment. Besides this, it costs DESY a lot of money and DESY may decide to spend this money for other purposes. We do not plan to remove the Siemens compiler in the near future, but you will be in a better position if you move to the IBM compiler as soon as possible. The obvious message after this introduction is:

*Do not use the Siemens compiler any longer; move to the IBM compiler !*

## How to move to the IBM VSFORTRAN Compiler

The general principle is simply to

- change a few things where Siemens and IBM behave differently
- recompile everything with the IBM compiler
- use the IBM VSFORTRAN version of all libraries

After these steps your program will run - at least in theory.

The differences mentioned above are mostly related to files, especially file names, file status and closing of files. For the Siemens compiler, the FILE parameter in the OPEN statement asked for the name of a *dataset*, but for IBM VSFORTRAN this would be interpreted as a *DDname*. To use a dataset name in the OPEN statement with the IBM compiler, you have to precede it with a slash ('/') as in the following example:

```
OPEN (UNIT=44, FILE='F99ABC.MY.DSNAME')
```

*for Siemens*

has to be changed to

```
OPEN (UNIT=44, FILE='/F99ABC.MY.DSNAME')
```

*for IBM*

More details under NEWLIB in **Help OLDFORT** in section

"How to move to the VS FORTRAN compiler"



## • Remote Login

*by Jan Hendrik Peters*

In your daily work you quite often have to work on different computer systems. This is not only true if you visit another institute and want to connect to your home computer in order to check mail. Even when you are located at DESY, it is sometimes necessary to connect to another machine to retrieve information, copy datasets, set up a graphics session, run a text formatter like TeX or something similar.

In this article I want to give an overview of the various possibilities of connecting to different computer systems. How this is done depends on the type of computer you are currently logged on and on the type of machine you want to connect to.

A connection between two machines can be set up if both are on the same network. The most common networks within the HEP community suited for remote login are DECNET (connecting VAX computers), the Internet running IP (IP = Internet Protocol), SNA-networks (SNA = System Network Architecture, connecting IBM computers) and X.25-networks running on public lines or private leased lines.

Each computer on these networks is identified by a unique name or number. Some computers are connected to more than one of these networks. Correspondingly, a computer may have different names on different networks, which sometimes can be quite confusing. Beside the networks already mentioned, which allow remote login, there are some mail-only networks like the EARN/BITNET (European Academic Research Network / Because Its There Network – originally connecting IBM computers) and the X.400 network of the DFN (DFN = Deutsches ForschungsNetz) which again assign different names to the computers, adding to the confusion.

Just to give you two examples from within DESY: the IBM computer is called DHHDESY3 on EARN, DSYIBM on DECNET, dsyibm.desy.de on Internet (mvs can be used as an alias for dsyibm), while the central VAX cluster is called DESYVAX or DHHDESY5 on EARN, VXDESY on DECNET,

and vxdesy.desy.de on Internet. At DESY we try to keep the convention that all computers that are connected to DECNET carry their DECNET name in front of the Internet domain name which is desy.de for all DESY computers and ifh.de for computers at Zeuthen. Unix computers are located in the same Internet domain and have their hostname attached to the domain (e.g. the ZEUS DEC5000 stations are called zow01..., zow03..., zow05.desy.de and so on, the RS6000 at Zeuthen are called r6l3.ifh.de (L3 machine) and r6rz.ifh.de (computer center)).

As a general rule of thumb: if you want to set up a connection between two computers of the same type you should use the proprietary network delivered with the machine, i.e. DECNET for VAXes, SNA for IBMs, and IP for Unix workstations. As for any rule, there are of course exceptions whenever the described way does not exist or using another network is much faster and more effective than the native one.

In the following sections I will describe what commands are available on which machines and how to use them. If you want to make full use of the remote computer you should make sure that the product you use has the correct terminal emulation built in. To work with a VAX you will need some kind of VT-emulation (VT = Virtual Terminal) say VT100, VT220 or the like, for IBMs a 3270-emulation is necessary. Some of the 3270-emulations have been written at DESY or modified in such a way that they even allow remote graphics sessions.

For people who already know all this but have just forgotten the commands or want to use a new method, I have included a small matrix of connection commands on the next page. Online helps for all these commands can be found on the corresponding machines.





Commands available for Remote Login				
from	to VAX	to IBM	to Unix	special tools
DESY VAX	set host telnet	set host/sna	telnet	set host/x29 set host to DECstations
DESY IBM	x telnet ((VT))	x telnet logon applid(samon)	x telnet	((PAD)) ((CERNVM))
DESY Unix	telnet rlogin	telnet 3270 tn3270	telnet rlogin	dlogin via DECNET from DECstations

## VAXes

The usual command to set up a remote connection from a VAX is `set host`. With this command you can start a session via DECNET to another VAX or to other computers running a DECNET connection software (e.g. Interlink on IBM, KiNET on Unix workstations), or via a X.25 line by giving the corresponding X.25 or dte-number. To use one or the other method, qualifiers have to be given like `/sna` or `/x29` when issuing the command. A few examples:

```
to VAX : set host zeus02
        set host vxcrna
to IBM : set host/sna dsyibm
        set host/sna cernvm
via X.25: set host/x.29 45050252310
        (IBM line mode)
```

Further parameters might have to be given when setting up the connection. Consult the following helps on the VAX:

- o `help set host`
- o `help sna set host`
- o `help P.S.I. set host`

With the `telnet` command you can set up a connection to a remote computer over IP. To do so you have to know the IP-number or IP-name of the remote machine. The IP-number consists of four parts separated by dots, the IP-name usually consists of the hostname and the internet domain name. Some examples:

```
telnet vxcrna.cern.ch
```

```
telnet nikhefh.nikhef.nl
```

```
telnet slacvm.slac.stanford.edu
```

With the IP-number it will look like this:

```
telnet 131.169.35.70 (our central VAX)
```

To connect to an IBM computer you need a telnet version with a 3270-emulation which unfortunately is not yet available on the central VAX. Typically the command would be called `tn3270` and would be used exactly as shown above. The IP-address for our IBM is 131.169.1.253 or `dsyibm.desy.de`, for our central VAX clusters it is 131.169.35.70 or `vxdesy.desy.de`.

## IBM

Several IBM computers in Europe are connected by an SNA-network. Computers on this network can be accessed from any IBM terminal at these sites without logging into the local computer. At DESY you can obtain a selection menu of all sites reachable, by typing

```
logon applid(samon)
```

or

```
m2 applid(samon)
```

instead of the usual `logon` or `m2` you type to access our local IBM computer. Once you have entered the SAMON (= Session Application MONitor) each of the nodes listed can be accessed by typing their nodename. Further information about the listed nodes can be obtained by typing their *node-name/I*. A similar service is available at CERN on the VM computer. Here you access the SNA nodes by typing `dial net` in the command line of the `logon`



# DESY Computing Newsletter No. 1



panel and selecting the node wanted (e.g. DSY for the DESY IBM). To disconnect from a remote VM host type `VMEXIT`, from a remote TSO host press PA1 and type `LOGOFF` after the logon prompt.

A more elegant and usually faster way is to set up your session on the remote host via `telnet`. The IBM telnet implementation sets up your session like on a VM system. This means you have to press the PA2 key (or sometimes the CLEAR key) if on the bottom right hand side of the screen the string `HOLDING` or `MORE...` appears. To start a telnet session under NEWLIB, type

```
x telnet cernvm.cern.ch
x telnet 140.181.64.11 (GSI IBM/MVS)
or the like. Pressing the PA1 key during the session will allow you to give a command to the telnet program (type ? for help).
```

With this implementation it is of little use connecting to a VAX or Unix station, although it is possible. For instance, the VAX fullscreen editors, Unix screen oriented outputs or the *vi* won't work under the 3270 emulation. For simple tasks, like checking your mail, it is certainly sufficient. With the same command as shown above you can connect to these types of computers.

At CERN the corresponding command on VM is:  
`tcpi.ibm telnet dsyibm.desy.de`  
or shorter  
`telnet dsyibm.desy.de`

With the PAD service (PAD = Package Assembly Deassembly) on the IBM you can set up a dialled X.25 line to other hosts. X.25 or dte-numbers (DTE = Date Terminal Equipment) on the WIN (= WissenschaftsNetz) in Germany start with 45050..., to dial to or from foreign countries a prefix has to be used (262 for Germany, 228 for Switzerland). If an IXI-line (IXI = International X.25 Interconnection) will to be used, 45050 has to be replaced by 02043623. All DESY computers can be reached with the extension 3540 and a two digit number assigned to the host: `desynet = 01`, `heracc = 02`, `IBM full screen = 08`, `VXDSYA = 80`, `VXDSYB = 81`, i.e. `VXDSYB` via WIN = 45050354081 and via

IXI = 02043623354081. To access the PAD under NEWLIB, type  
`((PAD))`

and after the PAD prompt `PAD>` type  
call 45050354080 (our VAX machine `vxdsya`)  
or any other number. With CERN we use a special mechanism that allows us to use a special number without the country prefix:

```
call 45050354030 (the CERNVM)
or even shorter with the NEWLIB clist
((CERNVM))
```

To leave the PAD after logoff from the remote system, type `EXIT`.

## Unix workstations

The natural choice on the Unix workstations is the `telnet` command. It will connect to other hosts via IP. To connect to an IBM you need a telnet version with a 3270 terminal emulation. Usually it is called `tn3270`, but at DESY you will also find `3270` which then is a DESY product or contains modifications allowing graphics sessions on the IBM host from your Unix workstation. Depending on the size of the window on your workstation the corresponding screen on the IBM will be set up when using the 3270 version of telnet. IBM screens at DESY have sizes of 80x24, 80x32, 80x43 and 80x48 (number of columns x number of lines). The syntax of the commands are as described above in the previous two sections:  
`tn3270 dsyibm.desy.de`  
`telnet vxdesy.desy.de`  
and so on.

Beside the telnet command, workstations allow a remote login with the `rlogin` command, the DECstations in addition have the `dlogin` command. `rlogin` will set up a session via IP, `dlogin` via DECNET. Valid commands are e.g.:  
`rlogin r0202` (the R2 HP station)  
`dlogin vxdesy`

I am quite sure that in this already long list I have missed out a few of the commands you already use and like. So please report anything that is missing, so they can be included in a forthcoming paper.



## • Usage of CERNLIB at DESY

by Harald Keller

All CERN program libraries will in future be offered at the same frequency as at CERN. CERNLIB releases are normally coupled with the publication of the CERN Computer Newsletter (or short CNL). Once the newsletter has been published, the latest version described in the newsletter will be released as the production version at CERN. Shortly after, the new libraries will be sent to all other sites that are on the mailing list. While the versions for the VAX and the workstations (IBM RS6000, DEC5000, APOLLO, SUN, Silicon Graphics) then can be used immediately at the remote sites, the version for the IBM has to be customized to fit into the MVS environment. Therefore this version needs some time before it is available.

Currently the version as of CNL200 is the latest version distributed by CERN. CNL201 was issued some time ago already, but the corresponding version of the CERN libraries are just now being distributed by CERN to the other labs. Thus we are installing the new versions just now, and some pieces have been installed already. The *.tar*-files for the UNIX-workstations are available for installation as well.

### Who is responsible

Harald Keller (userid R01KEL ext. 2501) is responsible for the installation of CERNLIBs on the IBM from now on. He will install all future versions on the IBM and he will provide all files for the different workstations (see below). Please direct suggestions and complaints to him.

### CERNLIBs on VAXes

On the central VAX cluster 2 versions of the CERNLIBs are currently available, both containing the complete distribution. Available are the versions corresponding to CNL200 and the preceding version. The older version is the default one, and the transition to CNL200 will be announced later.

Two command files are prepared to access the CERNLIB an VXDSYA and VXDSYB :

```
DISK$LIBRARY:[DESY]CERNLOGIN.COM
```

```
DISK$LIBRARY:[DESY]CERNSTART.COM
```

but most LOGIN procedures - especially for ZEUS - contain them already. The commandfiles for the CNL200 version are:

```
DISK$LIBRARY:[NCERN.PRO.MGR]CERNLOGIN.COM
```

```
DISK$LIBRARY:[NCERN.PRO.MGR]NCERNSTART.COM
```

More information can be found in

```
CERN_ROOT:[HLP]CERNLIB.HLP
```

### CERNLIBs on UNIX-workstations

CERNLIBs for all types of UNIX workstations are distributed in *tar*-format and they are accompanied by a *\$README* file for each machine.

All these *tar*- and *\$README* files are transferred to DESY as soon as they are available from CERN. All managers of workstations and workstation clusters, who want to install the CERNLIBs should take the DESY copy and install this copy, without extra traffic on the DESY-CERN link. We will always have the complete set of files available for all types of workstations. The filenames are of the form

```
R01KEL.CERN.CNL<ver>.$README.<sys>
```

for the *\$README* files and

```
R01KEL.CERN.CNL<ver>.TAR.<sys>.<lib>
```

for the different libraries and PAM files.

<ver> is the CNL version, e.g. 201

<sys> denotes the workstation type, e.g. A10

<lib> identifies the various libraries and PAMs

For the following systems CERNLIBs are available:

APO	Apollo 3000
A10	Apollo 10000
DEC	DECstation 5000
IRS	IBM RS6000
SGI	Silicon Graphics Iris
SUN	SUN SPARCstation



## CERN Load Libraries on IBM

There are two "lines" of CERN LOAD libraries at DESY: One line carries in its name the relevant CNL number, the other line is named along an OLD-PRO-NEW scheme, where

- OLD - the last but one production version
- PRO - the current or 'production' version
- NEW - the recently installed 'test' version

### Current state of CERNLIBs

CNL LINE	OLD-PRO-NEW Line
Every set of libraries corresponds to a CNL version	from the CNL Line 3 generations are maintained.
<currently not available> r01utl.cern.suffix *) r01utl.cern.cnl200.lib.suffix	r01utl.cern.OLD.suffix r01utl.cern.PRO.suffix r01utl.cern.NEW.suffix
and in addition from CNL 201 already	
r01utl.cern.cnl201.lib.graflib r01utl.cern.cnl201.lib.pawlib	r01utl.cern.NEW.graflib r01utl.cern.NEW.pawlib

\*) Version of November 1989

The explicit specification using the CNL-style names has the advantage of an unambiguous specification which stays valid forever, but does not keep up with new developments or updates. The *OLD-PRO-NEW* style allows the user to keep up with new versions automatically without ever changing his JCL.

Up to and including the version CNL200, DESY had supplied only KERNLIB, GENLIB and PACKLIB. Starting with CNL201, GRAFLIB and PAWLIB will be supplied as well, and these are already installed.

The contents of the various libraries is as in the following table.

"suffix"	containing
KERNLIB	KERNIBM KERNFOR KERNBIT KERNNUM
GENLIB	GENLIB
PACKLIB	KERNLIB <i>is included here to save one extra concatenation for PACKLIB users</i>  ZEBRA  containing: ZEBRA ZEMQ ZEFQ ZETEST DZEBRA FZCONV RZEBRA RZTEST
and from CNL201 on:	
GRAFLIB	HIGZ HPL0T5
PAWLIB	PAW COMIS SIGMA

The 'production' set of CERNLIBs with complete names is therefore:

- R01UTL.CERN.PRO.GENLIB
- R01UTL.CERN.PRO.KERNLIB
- R01UTL.CERN.PRO.PACKLIB

## Compilers / Dynamic COMMON

The *ONLY* FORTRAN compiler used for the CERN-libraries is IBM VSFORTRAN (FVS), so there are NO parallel versions any longer. In all compilations beginning with version 'CNL200', 'CNL201' or 'NEW' the user **MUST** use the compiler option

**DC(PAWC)**

because the packages that use PAWC are compiled



with this option as well. If you omit this option, your program will not run as expected and produce strange results (or abend) without any error message during compile or link.

There are no new CERNLIBs for the FORTRAN IV and Siemens compilers, but if you need an old version anyhow, take the libraries:

- R01UTL.CERN.GENLIB4
  - R01UTL.CERN.KERNLIB4
  - R01UTL.CERN.PACKLIB4
- or
- R01UTL.CERN.GENF77
  - R01UTL.CERN.KERNF77
  - R01UTL.CERN.PACKF77

## CERN PAM files

The pool of software sources for different machines and system programs is maintained on several files called PATCHY Master files, or PAM files for short. A PAM file is a sequential data set, packed into binary FORTRAN records for rapid handling, both by the program PATCHY and by the auxiliaries.

A PAM file is divided into named divisions called "*patches*", which are normally sub-divided into "*decks*". A "*deck*" corresponds to a routine. Those parts of source code included in a "*deck*" which match your machine and system program are controlled by flags.

Every PAM file on the IBM corresponds to a CNL version. You find the PAM file for the 'package' and the CNL 'version' you want under:

```
R01KEL.CERN.version.PAM.package
```

An example for PAM files - with complete names:

```
R01KEL.CERN.CNL200.PAM.KERNIBM
```

```
R01KEL.CERN.CNL200.PAM.KERNFOR
```

To obtain e.g. a list of all CNL200 PAM files please give the command:

```
LC R01KEL.CERN.CNL200.PAM
```

## Documentation

### *Printed Documents*

CERN program manuals will in future be imported from CERN through the DESY User Support Group (USG) only. An arrangement to this effect has been made with CERN. If you need any personal copies of CERN documentation please address the User Consulting Office (UCO) during UCO opening hours.

### *Online Documents*

The current state of the CERNLIBs on the IBM will be always available with

```
help CERNLIB
```

under NEWLIB.

Some (old) documentation is available online on the central IBM. For the CERN program library manual online ( -old- version of 25. July 1988) please see:

```
help CERNMAN
```

There is also an - old - IOPACK manual online:

```
help IOPACK
```

For informations extracted from the PAM files please see:

```
lds 'R01UTL.CERN.DOCUMENT'
```

and then list the member with the package name of your interest.



## ● FORTRAN Tips

by Michael Behrens

This is a small collection of FORTRAN-related tips that will be continued in the future. I would appreciate very much additional material from the readership.

### **BLOCK DATA** (*on IBM*)

There are currently 2 problems when using BLOCK DATA with the IBM VSFORTRAN compiler.

- special care is needed to include it when linking the program
- it compiles VERY slowly

All previous compilers on the IBM made it rather simple to include BLOCK DATA contained in a separate member when linking the program. An EXTERNAL statement specifying the name of the BLOCK DATA subprogram was sufficient.

With VSFORTRAN the semantic of the EXTERNAL statement has changed. The EXTERNAL statement merely specifies some properties of the names specified. The effect is, that names specified in an EXTERNAL declaration will be treated properly if they are used somewhere in the program unit containing the EXTERNAL statement. Unfortunately, the name of a BLOCK DATA subprogram is never used anywhere in the program, and it will therefore never be included in the resulting program. There are no error message at compiletime or runtime, but the running program will produce strange results.

Special commands for the Linkage Editor will help, but they are somehow cumbersome and most people have no experience with advanced Linkage Editor functions. To do this the procedure VFORT-CLG or any other ...CLG procedure must be used with

```
//L.SYSIN DD *  
  INCLUDE SYSLIB( block-data-name)
```

This method has to be used with some standard MC

generators where the BLOCK DATA section is part of the MC load library.

A simpler method is to place the BLOCK DATA in the same member as a dummy procedure that is called once during an initialization phase, e.g. from the main program. As the BLOCK DATA is now contained in the same textual unit as the additional dummy procedure, it will be automatically included with it in the resulting program.

Several users have noticed that BLOCK DATA compiles very slowly on the IBM. Compilation speed is up to 15 times slower for BLOCK DATA than with the Siemens compiler. This problem is currently being discussed with IBM but one cannot expect a solution in the near future. Often BLOCK DATA is included with the main program for reasons mentioned above, and is therefore compiled again and again. Here the above method of using a dummy procedure will help as well, because the BLOCK DATA will be compiled only when really needed, and the excessive compilation time will less often occur.

### **Warning for uninitialized variables** (*on IBM*)

The current release of IBM VSFORTRAN will produce warnings at compile time for (potentially) uninitialized variables. These warnings require OPT(2) or OPT(3), as these warnings are a byproduct of program flow analysis, which is done only with the higher optimization levels. Under NEWLIB you simply call the compiler by

```
C0mpile <member> CPRM OPT(2) or OPT(3)
```

### **REAL vs. DOUBLE PRECISION** (**REAL\*4 and REAL\*8**)

A common, but wrong assumption is that operations on DOUBLE PRECISION ( or REAL\*8) data are slower than operations on REAL\*4 data. This assumption was true in former times, when operations were performed on 32 or even less bits at a time, but it is no longer true in general.



It is still true for the VAX, where REAL\*8 takes considerably longer than REAL\*4, but it is not true many other machines.

The larger IBM machines, including the DESY machine operate on 64 bits anyway. So no significant penalty in terms of processor time is introduced when moving from single to double precision on the IBM.

PC's normally use a coprocessor to operate on floating point data. These coprocessors work on DOUBLE PRECISION data anyway and there will be no overhead involved in the computation. PC's without coprocessor are of course slower, but the speed for both floating point formats will still be the same.

On workstations things are progressing rapidly. In many cases REAL\*8 or DOUBLE PRECISION is still slower, but a trend towards 64-bit machines will lead to almost the same speed for REAL and DOUBLE PRECISION in the future.

## INTEGER\*2 (on IBM)

Short integer arithmetic on the IBM is effectively slower than normal integer arithmetic. In addition to that, the vector unit is unable to operate on INTEGER\*2 data. The final result is that code with INTEGER\*2 data will not vectorize on the IBM.

## • Fortran 90 - the current state

*by Michael Behrens*

The new FORTRAN standard, usually called *FORTTRAN 8X* during the last years, has got a new name *Fortran 90* and its long lasting development has now come to an end. The draft for the new international standard was already completed last summer and the remainder of last year had been used to do the final editing of the new revised standard.

After public review the final ballot by the International Standards Organization( ISO) was completed in February 1991. The results were

- 19 countries : yes
- 1 country : no (Japan)
- 1 country : abstention (Thailand)

The japanese vote might eventually change to YES if certain edits are performed, but a single dissenting vote will do no harm to the standardization process. This ballot was the last vote on the technical contents of the revised language. Quite a few bureaucratic steps on several levels of ISO are still to be done, but this will not affect the contents of the language at all.

So the new standard is effectively there. Several vendors have been working on the development of compilers for the new language for quite some time now, and I would expect the first compilers to appear within a year. These early compilers might not implement all of the very advanced functions of the new language immediately, but they could well serve as a training ground.

Several talks presented the (emerging) new language at DESY during the past years. There will certainly be updates this summer at DESY. In the meantime, interested people can find some literature in the DESY library.

- M. Metcalf, J.Reid : Fortran 90 explained
- W. Brainerd, C.Goldberg, J. Adams : Programmers Guide to Fortran 90



The first book concentrates more on the new things in Fortran 90, the second starts almost from scratch and does not refer to previous FORTRAN versions.

## • **The Database System DB2 at DESY**

*by Peter Dobberstein*

### **General Remarks about DB2**

DB2 is a relational database system running under the operating system IBM/MVS. The corresponding IBM/VM product is called SQL/DS. Both of them are accessible interactively under the common user interface QMF (= Query Management Facility).

One can perform queries on a database using several languages: Interactively within QMF, one can use

- The QMF commands,
- the QBE keywords (Query By Example),
- the Prompted Query commands,
- the SQL statements (Structured Query Language),

depending on the profile and the submenu chosen (see *QMF Literature* below). QMF also supports the generation of data base reports.

In a higher programming language (e.g. FORTRAN, C), one is able to include SQL statements. The programs can be run interactively and in batch.

*WE RECOMMEND THAT THE BEGINNER BECOME FAMILIAR WITH SQL FIRST.*

*This language can be used interactively and in batch, as well as within the framework of many other database systems.*

To support the usage of QMF and DB2, the reference summaries of these products are available at the User Consulting Office. For more detailed information ask

Klaus Dahlmann (R01DAH, ext. 3622)  
or Peter Dobberstein (R00DOB, ext. 2732),

and/or consult the literature listed below.





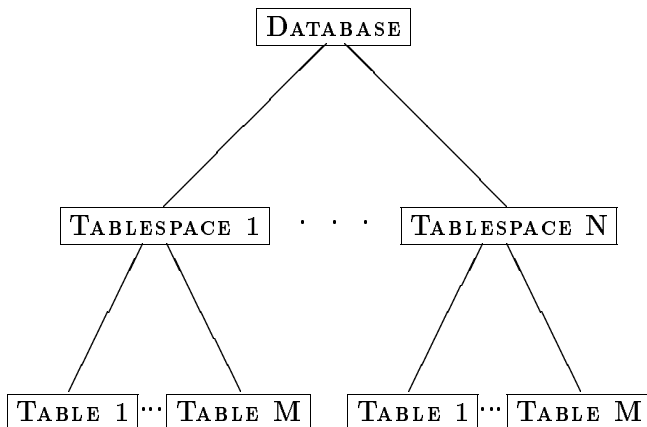
## Access to DB2 at DESY

To gain access to the DB2 system, please contact K. Dahlmann (ext. 3622). He will establish a DATABASE for you. At the time of the initialization, you have to specify:

1. The UserID the DATABASE belongs to,
2. the maximum disk space required for your purposes (in MBytes),
3. whether your DATABASE requires a regular backup or not.

Within your given resources, you are the DATABASE administrator and are authorized (for example) to create and to drop TABLESPACES and TABLES.

The terms DATABASE, TABLESPACE and TABLE are correlated as shown in the following diagram:



In words: A DATABASE contains several TABLESPACES and each TABLESPACE consists of several TABLES, which contain the data.

## DB2 Interactive Usage (QMF) and the DESY Phone Book

Interactive access to DB2 via QMF is realized using IBM's ISPF environment. To invoke QMF, the NEWLIB clist ((QMF)) should be used. Typing in

((QMF))

allocates the required files and builds up the ISPF environment. In the first menu, one can choose the print destination for QMF (which remains valid until it is redefined), enter QMF, or use the DESY phone book (which is organized as a DB2 table) as a first public application.

## DB2 Usage from High Level Languages

It is possible to access DB2 data from a high level programming language (e.g. FORTRAN, C). To achieve this, the following steps are required:

1. Embed SQL statements defining your DB2 access into the host language (FORTRAN, C, ...).
2. Precompile the source code. The precompiler resolves the SQL statements into host language elements and creates a so-called database request module which is used at run time.
3. Compile and link-edit the output of the precompiler as usual.
4. Run the program using a special environment which connects to the database. This environment is built up by a command processor named DSN, which has to be started from batch-TSO. DSN uses the database request module generated by the precompiler to create an 'application plan'. The 'application plan' contains information about the access path to the data and is used together with the application program.

To relieve the user of complicated JCL problems, the computing centre supplies some procedures which are similar to the well known 'COMPILE', 'LINK', and 'GO' procedures. At present these procedures are available for the VS-FORTRAN and IBM's C/370 compilers.

The equivalence between the existing VFORT-procedures (pure VS-FORTRAN) and the new DB2F-procedures can be seen in the following table:



	VS-Fortran	VS-Fortran with embedded SQL
Compile	VFORTC	DB2FC
Compile, link	VFORTCL	DB2FCL
Compile, link, go	VFORTCLG	DB2FCLD
Go	VFORTG	DB2FD

The corresponding C/370-procedures are equivalenced as follows:

	C/370	C/370 with embedded SQL
Compile	C370C	DB2CC
Compile, link	C370CL	DB2CCL
Compile, link, go	C370CLG	DB2CCLD
Go	C370G	DB2CD

Look into the procedures which are stored in 'SYS1.DESYPROC' for additional information.

```

Before running one of these procedures, one has
to ensure the existence of the following data
sets:
- 'userid.SRCLIB.DATA'
  (DCB model = 'R01DCB.CARDIM')
- 'userid.DBRMLIB.DATA'
  (DCB model = 'R01DCB.CARDIM')
- 'userid.RUNLIB.LOAD'
  (DCB model = 'R01DCB.LOAD')

Otherwise, the procedures will produce unpre-
dictable results.

```

## DB2 Usage from other Systems or Computers

There are a lot of connections from other systems, databases or computers to a DB2 database. On the IBM site, interconnections from a DB2 database to SQL/DS (VM), AS (Application System) and other DB2 databases are possible. Furthermore, DB2 tables can be made readable from ORACLE and VAX Rdb/VMS.

At present, the interface to AS is implemented and tested. Other desired connections are not implemented yet, but are being explored.

## DB2 Literature

The following is a list of IBM literature which is of interest to the end user. All these books are available at the User Consulting Office. The reference summary is being freely distributed.

- DB2: Application Programming and SQL Guide, SC26-4377**
  - Using Interactive SQL
  - Coding SQL in Your Application Program
  - Testing and Executing Your Application
  - Programming for Special Purpose Interfaces
- DB2: SQL Reference, SC26-4380**
  - Concepts
  - Language Elements
  - Functions
  - Queries
  - Statements
- DB2: Command and Utility Reference, SC26-4378**
  - Commands
  - Utilities
- DB2: Reference Summary, SX26-3771**
  - SQL Reference Summary
  - Command and Utility Reference Summary
- DB2: Messages and Codes, SC26-4379**
  - SQL Return Codes
  - DB2 Messages
  - DB2 Codes

## QMF Literature

The following is a list of IBM literature which is of interest for the end user. All these books are available at the User Consulting Office. The reference summary is being freely distributed.

- QMF: Learner's Guide, SC26-4231**
  - Contains step-by-step instructions on how to use QMF.
- QMF: Advanced User's Guide, SC26-4343**
  - Provides guidance for using more advanced QMF functions that the novice may not require.



3. **QMF: Reference, SC26-4344**
  - Contains information on QMF commands, SQL statements, and the use of QMF form panels to format reports and charts.
4. **QMF: Reference Summary, SX26-3752**
  - Provides a quick reference summary of QMF commands, SQL statements, form panels, QMF symbols, data types, and system tables.
5. **QMF: Query-by-Example Guide and Reference, SC26-4364**
  - Explains how to write queries in the Query-by-Example language and describes QBE keywords.
6. **QMF: Application Development Guide for MVS, SC26-4237**
  - Discusses the use of QMF's application support facilities for MVS.

## • Questions and Answers from the UCO

*by Katherine Wipf*

The User Consulting Office answers quite a variety of questions, but usually only a few people hear the answers. In the hope of bringing the answers to more frequently asked questions to a larger audience, I have collected a few questions from our log file.

### . . . about FORTRAN

*Question:* How can I get my program to wait or a certain number of seconds before continuing execution?

*Answer:* There is a routine in the DESYLIB called DELAY, which will cause a program to hold execution for a given number of 26 microsecond time units. Call the routine as follows:

```
CALL DELAY(NHSEC)
```

where NHSEC is the number of 26 mikrosecond time units the program should wait before continuing.

---

*Question:* I'm compiling and running my FORTRAN program in a batch job using VFORTCG. How can I prevent my program source from being listed in the job output?

*Answer:* You can use the compiler option NOSOURCE to suppress your source listing. You would select it in your JCL as follows:

```
// EXEC VFORTCG, CPRM=NOSOURCE
```

---

*Question:* When I try to link my program, I get a message that tells me NO ENTRY POINT ASSIGNED. What does this mean?

*Answer:* When the name of your program is different from the name of the member which contains it, it is necessary to give an entry point in the link panel. In the MEMBER NAME => field you should enter the name of the member and in the ENTRY POINT => field the name of the program.



*Question:* My program is producing so many UNDERFLOW EXCEPTION messages that my job exceeds the line limit for class A. Is there any way of suppressing these messages?

*Answer:* There is a Fortran subroutine called XUFLOW which takes one argument. It is called as follows:

```
CALL XUFLOW(k)
  where k: 0 - program interrupts and
            error messages are sup-
            pressed if an under-
            flow occurs
          1 - program interrupts caused
            by exponent underflow
            occur
```

NOTE: These UNDERFLOW EXCEPTION messages are generally an indication that your program contains mathematical errors (i.e. - inadvertently dividing by a very large number, or possibly incorrect parameter handling). As a general rule it is a better idea to locate and remove the error than to suppress the error message.

## . . . about MAIL

*Question:* When I try to call SENDMAIL or READMAIL from the ((MAIL)) clist, I get the message:

```
****ERROR**** IN GETMAIN
NOT ENOUGH SPACE AVAILABLE
What does it mean?
```

*Answer:* It means your logon size is too small. The next time you logon, look at the panel where you enter your password. There is field marked SIZE ==> which should be left blank. This will always give you the default (maximum) logon size.

*Question:* How can I include a NEWLIB member in a message I'm sending with the ((SENDMAIL)) clist?

*Answer:* You need to use a parameter when you call SENDMAIL. This is the correct syntax:  
((SENDMAIL)) include *membername*

*Question:* I often send mail to someone with a complicated electronic mail address and would like to create a nickname for it. The problem is that I'm not quite sure how to divide it up into node, network, etc.

*Answer:* There is a new option for the nicknames file which lets you enter the entire address rather than splitting it into node, gateway, etc. You need to enter the following in your nicknames file:

```
:nick.nickname
:email.userid%node@othernode
```

*Question:* I want to send mail to the node HERAMM, but since it isn't on BITnet, I can't send to *userid@HERAMM*. How do I reach this node?

*Answer:* Since the HERAMM is a DECnet node, it can be reached by using the DESYVAX as a gateway. Give the address as follows:  
*userid%HERAMM@DESYVAX*

## . . . about other Topics

*Question:* I'm trying to use the VAX tape machine near the external printers, but I can't put the device ON LINE. What am I doing wrong?

*Answer:* There are a few things one has to remember when using the VAX tape machine:

1. After placing the tape in the machine, always remember to press down the ring which you will find at the center of the tape.
2. Press the LOAD REWIND button.
3. Once the tape has stopped moving, press the ON LINE button.

*Question:* I am using interactive GEP in the one terminal mode and the X command doesn't seem to be working. When I give the X command my screen remains blank, but if I use SEL and D my picture is drawn without any problems. Why does this happen?



*Answer:* This is a bug in the GEP program which occurs when GEP and the graphics package used to display GEP pictures give `<CLEAR SCREEN>` commands in the wrong order. If you follow your X command with an RX command, your picture will be drawn.

---

*Question:* I'm not getting messages about my jobs finishing or new mail arriving until the next time I log on. How can I fix this?

*Answer:* This effect is caused by defective TSO profile variables. There is a command for re-establishing the variables. All you have to do is enter the following command in the command line:  
`INTERCOM`