BNL 50967

MASTER

**BROOKHAVEN NATIONAL LABORATORY
COMPUTER USER'S GUIDE**

PERMANENT FILE GUIDE

THOMAS S. CLIFFORD

**CENTRAL SCIENTIFIC COMPUTING FACILITY
APPLIED MATHEMATICS DEPARTMENT**

**BROOKHAVEN NATIONAL LABORATORY
ASSOCIATED UNIVERSITIES, INC.**

**UNDER CONTRACT NO. EY-76-C-02-0016 WITH THE
UNITED STATES DEPARTMENT OF ENERGY**

THE DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

BROOKHAVEN NATIONAL LABORATORY COMPUTER USER'S GUIDE

PERMANENT FILE GUIDE

THOMAS S. CLIFFORD

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

**CENTRAL SCIENTIFIC COMPUTING FACILITY
APPLIED MATHEMATICS DEPARTMENT**

**BROOKHAVEN NATIONAL LABORATORY
UPTON, NEW YORK 11973**

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

EB

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy (DOE), nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use would not infringe privately owned rights.

Printed in the United States of America

Available from

National Technical Information Service

U.S. Department of Commerce

5285 Port Royal Road

Springfield, VA 22161

Price: Printed Copy \$4.50; Microfiche \$3.00

May 1979

350 copies

	<u>Page</u>
1. Introduction	1
2. Basic Terms and Functions	3
2.1 lfn,pfn,ID,cycle,permissions	3
2.2 Creation of PF's CATALOG	5
2.3 Use - ATTACH	7
2.4 Annihilation of Files - PURGE.	8
2.5 Passwords, Permission, Multiread - Some Examples	8
2.6 Information - AUDIT.	10
3. CSCF File Management and Accounting	13
3.1 7600	13
3.1.1 Allocation Scheme	13
3.1.2 Allocation Algorithm.	14
3.1.3 Removing Files.	17
3.1.4 Backup.	17
3.1.5. Accounting.	18
3.2 6600	19
3.2.1 MFA	19
3.2.1.1 Archiving.	19
3.2.1.2 Removing Files From the System	21
3.2.1.3 File Backup.	22
3.2.1.4 Accounting	23
3.2.2 MFB	23
3.2.2.1 File Management.	23
3.2.2.2 File Backup.	23
4. User Backup	25
4.1 DUMPF.	25
4.2 LOADPF	26
4.3 Some Good Advise and Sample Jobs	27
5. The Use of Permanent Files Between Mainframes	29
5.1 SAVEPF, GETPF and PURGE.	29
5.2 Some examples	30
6. Private Packs	35

1. INTRODUCTION

An important measure of today's large computing systems is the amount of mass storage associated with the system. More important to the end user of today's computers than just availability of storage to the system is the ability to store files reliably and have them available for batch or interactive processing. On any system the amount of file space is finite and there is a great deal of pressure from many sources for the use of this file space, making its management extremely important and difficult. When disks become full, files must be removed, jobs abort, time is lost, people are unhappy.

The purpose of this volume of the CSCF Users Guide is to document the use and management of the file systems associated with the CSCF mainframes. The basic tools for storing, using, and discarding files are documented as well as scribed in terms of job control scribed in terms of job control statements but the user should be aware that there are FORTRAN subroutines and COMPASS macros to perform the same functions provided by the control statements.

Some material in this volume has been taken from documents filed in the AMD document file and from articles written for the CSCF Newsletter. The author is indebted to the authors of those articles as well as to those who have helped in the final drafts of this document.

THIS PAGE
WAS INTENTIONALLY
LEFT BLANK

2. BASIC TERMS AND FUNCTIONS

There are four types of files which can be associated with a user job: (1) queued files, which are associated with some I/O device; (2) tape files, which reside or will reside on a magnetic tape; (3) permanent files, which are known to the system and remain in the system after job termination; and (4) scratch files, which disappear at job termination. The subjects of this manual are the permanent files, often referred to as "perm files," "PF's," or "catalog files."

2.1 lfn,pfn,ID,cycle,permissions

Each of the files associated with a user job has a symbolic name, the Logical File Name, (lfn) by which they are referenced. An lfn is 1 to 7 alphanumeric characters long with the first character being alpha. Files with lfn of OUTPUT, PUNCH, PUNCHB, and P80L normally become queued files. Jobs which have entered a machine and which have not yet begun execution are also queued files (members of the input queue).

A permanent file is a mass storage file, known to the system permanent file directory, which can become immediately available to batch or interactive jobs. Permanent files are not destroyed at the end of a job and, at creation time, the user can establish protection for the file against unauthorized access or destruction by other users.

Permanent files are known to the system by a triplet (permanent file name, ID,cycle number) in the permanent file directory. When a PF is associated with

a job it must have an lfn, but that name is not part of the file name stored in the permanent file directory and thus may be different for different jobs.

The permanent file name (pfn) is one to forty alphanumeric characters with the first character being alpha. The ID field is used to insure uniqueness in PF specification. The ID can be one to nine alphanumeric characters. The common and suggested use of the (pfn, ID) pair is for each user to assign one ID to all his files and let the pfn describe the file. Having all a user's files under one ID is convenient because some of the the PF management tools key on ID. INTERCOM users sometimes confuse their user name, password, or two-letter ID with the permanent file ID's. There is no relation.

Up to five different files may be identified by the same (pfn,ID) pair with the cycle number (CY) used to distinguish the files. Even though there can be only five cycles (files) for each (pfn,ID), cycle numbers can range from 1 to 999. The normal use of cycle number is to distinguish different versions of the same file.

The permanent file directory contains accounting information, usage information, file size, and passwords as well as the (pfn, ID, cy). Passwords, defined when the file is placed in the permanent file directory, are used to regulate access to and control of PF's. There are four different types of permission allowed to a job which is using a PF:

read permission - job may read the file;

control permission - job may make a new cycle of the (pfn,id) or remove
the file from the permanent file directory;

modify permission - job may write on the PF;

extend permission - job may write at the end of the PF and thus lengthen the file.

Passwords are associated with (pfn,id) so that every cycle of a (pfn,id) has the same passwords defined.

It is possible for several concurrent jobs to be granted read access (even while one user has modify or extend permission). It is also possible to obtain exclusive access to a file. There are two types of read permission, exclusive read and multiread (MR). If a job has exclusive read for a file, no other job can be using that same PF at the same time. More than one job may be reading a file simultaneously if each job has only multiread permission for the file.

2.2. Creation of PF's - CATALOG

The control verb CATALOG is used to place a file in the permanent file directory and thus make it a permanent file. The details of all the options of the CATALOG control statement can be found in the SCOPE 3.4 and SCOPE 2 Reference Manuals. What is documented here is the most commonly used and suggested subset of all the options.

CATALOG(lfn,pfn,ID=yourid,AC=youraccount,RP=n,XR=pw,PW=controlpw,CY=m)

- lfn - The logical file name of the file to be placed in the permanent file directory. The named file must exist, be on a PF device (see below) and not already be a PF or queued file. The lfn itself is not entered in the PF directory.
- pfn - The permanent file name of the file being catalogued. If the user wishes the pfn to be the same as the current lfn of the file, the "pfn," may be omitted from the catalog card.
- ID=yourid - The PF ID must be specified. The ID is used along with the pfn to identify the file in the system. To help manage your files it is suggested that each user specify one ID for all his PF's.
- AC=youraccount - An account number must be assigned to each PF. On MFZ the PF account is the same as that of the job doing the CATALOG and the AC= is a null parameter. On MFA, the AC=youraccount will override the default account which is the account number of the job doing the CATALOG.

RP=n

- RP is retention period expressed in days. The RP is a value from 0 to 999, 999 being infinite retention. The RP parameter on the CATALOG statement is optional. A file whose retention period has elapsed will be removed from the system by the EXPIRE job. As you described in Section 3, files are removed from the PF directory for reasons other than their expiring. Retention period is an upper bound, a file will not remain longer than its RP; it may not last that long. Files with RP=0 may be expired any time during the day. RP=1 will be expired in the early morning of the day after the files were catalogued.

The default RP on MFA is 5 days

The default RP on MFZ is 1 days

XR=pw

- The XR parameter is optional; it is used to define all the passwords Except the Read (XR) password. XR=ABC on the catalog defines the control, extend, and modify password to be ABC. A job which is going to write on this PF or remove it from the catalog must specify PW=ABC.

PW=controlpw

- The PW parameter is optional; it is needed on the catalog to gain control permission only if a control password was defined in cataloging the first cycle.

CY=n

- CY is an optional parameter specifying a cycle number. It is generally not used. If CY is not used the cycle number of the file will be 1 on the first catalog of the (pfn,id) pair. If there is already a file catalogued with the (pfn,id) pair, then the new file will have a cycle number that is one greater than the highest catalog cycle.

As mentioned above, a file must exist on a PF device before it can be catalogued. The REQUEST control statement is used to signal to the system to build a file on PF device.

REQUEST,lfn,*PF.

The request statement is issued before the file is created (opened).

For files cataloged on the MFA system, a parameter has been provided to indicate to the system that the file should not be archived. The NA=1 on the catalog statement sets the no archive flag. (See Section 3.2.1.7)

2.3 Use - ATTACH

The control statement ATTACH is used to make a permanent file available to a job. Once a file is available to a job, it is referred to by its logical file name, the lfn.

ATTACH(lfn,pfn,ID=fileid, cycle indicator)

lfn	The lfn is the name the file will have for the job in which the ATTACH is done. It need not be the lfn used in the CATALOG.
pfn	The pfn is the permanent file name of the file. The lfn, can be omitted from the ATTACH statement. In this case the lfn of the file is the first seven characters of the required pfn parameter.
ID=fileid	The ID of the file must be specified unless the ID is PUBLIC.
Cycle indicator	If no cycle is indicated the highest cycle is attached. CY=n requests the specified cycle number be attached. LC=1 requests the lowest existing cycle number be attached.

If the file has been catalogued with an "XR=password" and the file is attached by concurrent jobs without "PW=password", then all of the jobs will be granted read and permission (multiread).

On MFA and MFB if a file has been catalogued without XR then all permissions are granted upon an attach shown above (no multiread).

On MFZ a file catalogued without XR and attached as shown results in multiread access to the file. See the examples below.

For most uses, it is suggested that PF's be read only files. If you wish to change a PF, write a new cycle of the existing file, then if the new cycle is correct get rid of the old cycle if you must and move on. This procedure provides a backup, the old cycle, and is safer and easier than writing on a PF. If you have a reason for writing on a PF, read the reference manual section on EX-TEND and ALTER and be careful.

2.4 Annihilation of Files - PURGE

PURGE is the control verb used to remove files from the PF directory. The PURGE control statement can be used on an attached file or on an unattached file.

PURGE(lfn)

lfn The lfn is the logical file name of an already attached PF. This statement will remove a file from the directory only if the file is already attached with control permission. If XR was used on the catalog of the file then the file must be attached with "PW=controlpw" to gain control permission for the file. See the examples in the next section.

PURGE(lfn,pfn,ID=fileid, cycle indicator,PW=pw,MR=0)

The file is attached on a purge (unless it is an archived file) and is thenceforth treated as a scratch file with the specified lfn. ID,pfn and cycle indicator are as in ATTACH.

PW=pw The control password, if defined, must be specified in order to gain permission to purge the file.

MR=0 For a file catalogued on MFZ, the user must specify "no multiread" in order to have control permission if no control password was defined for the file.

2.5 Passwords, Permission, Multiread - Some Examples

The following examples point out the basic use of CATALOG, PURGE and ATTACH. The differences between 7600 and 6600 should be noted.

2.5.1 7600

CATALOG(TAPE7,DATAFILE,ID=USG)

TAPE7 is catalogued with (pfn,id) of (DATAFILE,USG).
No passwords are associated with the file.

Retention period is 0 days; file will be gone tomorrow.

ATTACH(TAPE1,DATAFILE,ID=USG)

File catalogued above is attached as TAPE1 with read only permission; the MR=0 would be needed for write or control (PURGE) permission. Another job can attach the file in the same way while it is attached to this job.

PURGE(DATAFILE, ID=USG, MR=0)

(DATAFILE, USG) is removed from the PF directory. File is attached with an lfn of DATAFIL and remains local to the job.

* * * * *

CATALOG(NEWPL, BNLUPL, ID=USG, XR=FLEP, RP=999, PW=FLEP)

NEWPL is catalogued (BNLUPL, USG) with an infinite retention period. The control password is FLEP.

ATTACH(OLDPL, BNLUPL, ID=USG)

(BNLUPL, USG) is attached with lfn of OLDPL, read only permission.

PURGE(BNLUPL, ID=USG, PW=FLEP, LC=1)

The lowest cycle of (BNLUPL, USG) is removed from the PF directory.

* * * * *

2.5.2 6600

CATALOG(S1, JOB1, ID=USG, AC=1234)

File S1 is catalogued (JOB1, USG) with account number 1234. The retention period is 5 days.

ATTACH(JOB1, ID=USG)

Permfile (JOB1, USG) is attached with all permissions under an lfn of JOB1. No other user may attach this file while it is attached to this job.

PURGE(F1, JOB1, ID=USG)

Permfile (JOB1, USG) highest cycle is removed from the directory. The file is local to the job with an lfn of F1.

* * * * *

CATALOG(SOURCE, ID=USG)

File SOURCE is catalogued with pfn of SOURCE and ID of USG. File has the account number of the job doing the CATALOG. It will have default retention period of 5 days.

ATTACH(F1,SOURCE,ID=USG)
PURGE(F1)

FILE (SOURCE,USG) is attached with an lfn of F1 with all permissions and then purged.

* * * * *

CATALOG(NEWLIB,BNLU,ID=PUBLIC,PW=----,RP=999,AC=---)

A new cycle of BNLU is catalogued with ID of PUBLIC. The old XR password must be specified. Infinite retention is requested.

ATTACH(BNLU)

File (BNLU,PUBLIC) is attached with read only permission. Many users may have this file attached at the same time with read only permission.

2.6 Information - AUDIT

The utility program AUDIT provides the user with a tool for obtaining information on PF's in the permanent file directory. All the features of AUDIT and a description of the output are described in the SCOPE 2 and SCOPE 3.4 Reference Manuals. We will point out only the commonly used features.

AUDIT is a control statement which may be executed with the following options:

ID=yourid	Limits the AUDIT to files with specified ID.
AI=P	Produces one line of output per file. Without this parameter a more detailed two lines per file is produced.
LF=lfm	The lfm is the name of the file on which the output will appear. The default is OUTPUT.
AC=n	Limits the AUDIT to files catalogued with specified account. <u>AUDIT by account is costly, slow, and should be avoided.</u>
MO=R	6600 only. This option will limit the user to audit archived files.

The normal use of AUDIT would be

AUDIT(ID=yourid), or

AUDIT(ID=yourid, AI=P).

If you have been using the same ID on your catalog then the AUDIT by ID is all you should ever need.

MFZ users might find the PFSPACE described in Section 3 more useful for looking at their files than AUDIT but PFSPACE is slow and costly.

THIS PAGE
WAS INTENTIONALLY
LEFT BLANK

3. CSCF FILE MANAGEMENT AND ACCOUNTING

Trying to maintain the integrity of the PF file base and manage the always limited amount of space is a difficult and continuing task. The "system" does not do everything by itself; there are operational procedures followed to do daily and weekly dumps which provide backup in case of some file-damaging system failure. Each day the expired and illegal files must be removed from the system and measures must be taken to be sure that the PF's do not fill up all the disk space.

3.1 7600

3.1.1 Allocation Scheme

The total mass storage space devoted to PF's on MFZ is divided into n portions, with each portion assigned to a user group. A user group is defined by a set of problem numbers belonging to that group. A list of the active problem numbers and the groups to which they belong along with the allocated PF space is kept as a file on MFZ and can be listed with the following control cards:

```
ATTACH(PLIST,ID=SYSTEM)
COPY(PLIST,OUTPUT)
```

The space assigned to each group is apportioned (logically) into guaranteed space (g-space) and remainder space (r-space). The ratio of g-space to r-space is decided upon by the users in a particular group, but once decided

upon it is fixed. Files are allocated to guaranteed space by using an ID string with a prefix of ZZG on the CATALOG of the file. The ID string with a prefix of ZZR directs files to the remainder space while all other files fall into the peak usage space shared by all users.

3.1.2 Allocation Algorithm

Operations runs the PFALLOC job at least once a day to pare the PF base down to a size which will allow the system to run. Too much disk space allocated to permanent files reduces the space available for system use, such as scratch files and I/O files, and can cause the system to run out of operating space. The PFALLOC job removes files from the system according to the following algorithm:

- i For each group, check to see if the total space occupied by ZZG files is more than the allocated g-space for the group; if so, remove the last files catalogued until space occupied by the ZZG files for the group is less than the g-space allotted to the group.
- ii For each group, see if the total space occupied by ZZG and ZZR files is more than the total space allocated to the group; if so, remove (logically) the ZZR file which has been inactive (not attached) for the longest time. Repeat this logical removal until the total spaces occupied by ZZG + ZZR files is below the allocated space. Now delete from the list of files to be removed, starting with the smallest file, all the files which will fit into the allocated space. Only after all the smaller files which will fit into the allocated space have been deleted from the list of files to be removed, are files actually removed. This complex procedure is intended to try to prevent one new large file from pushing itself and lots of old little files out of r-space.
- iii If the total space occupied by non-ZZR and ZZG files is more than peak space, remove files from peak space until they all fit, starting with the files which have been inactive (not attached) for the longest time.

Some important observations follow from the above algorithm and its implementation. A file that overflows a group's guaranteed space will last less than one day. Be careful not to exceed your guaranteed space and use the pro-

gram PFSPACE described below to audit your g-space. Because the time of a CATALOG is not stored along with the date in the PF directory, it is possible to have a file removed from g-space because someone else overflowed the space the same day on which the file was "safely catalogued." Files can be pushed out of g-space if a file already in the space is made larger by a user extending his PF. Discounting the two unusual cases mentioned above, once a file is in guaranteed space it will remain, even if not used.

Files in reserved space are not safe from new files being catalogued in the r-space and pushing the "old" files out. One large file put into an already full r-space will push out many smaller files if they were not attached after the large file was catalogued. When files are removed from the system because they overflowed the group's available space the files are dumped to a disk pack and are kept there for one to three days. While the CSCF does not routinely restore files from these packs the files can be restored in case of some emergency.

Peak space is not very safe. Users dump large data files in this space not expecting it to last and indeed it does not. The program PFSPACE is available to allow users to check the current status of the storage space assigned to a group.

PFSPACE is a utility program to enable users to see how much of their allotted space is occupied, and by which files; this is accomplished by scanning the output of the AUDIT utility. Input consists of a group ordinal or name, or a problem number. Output consists of a list of files belonging to the group or problem number including, for each file, the creation, expiration and last access dates, and the size of the file, expressed in words, on the 7600. A summary shows how much of each type of space is in use, and what the group limits are for g-space and r-space.

Input data cards, one or more per request. FORMAT is (A1,I4,1A10).

Col. 1: G for group ordinal. P or + for problem number.

+ indicates that more data cards follow for this request.

Cols. 2-5: A group ordinal or problem number, in I4 format.

Cols. 6-25: Optional group name (when Col. 1 is "G" and Cols. 2-5 are blank).

DECK STRUCTURE:

```
jobname,STMFZ,TP0.
ACCOUNT(yourname,num)
AUDIT(LF=TAPE1)
ATTACH,TAPE13,GPLIST.
ATTACH,BNLU.
LIBRARY,BNLU.
MAP(OFF)
```

```
PFSPACE.
(7/8/9 end-of-record)
data deck
(6/7/8/9 end-of-job)
```

EXAMPLE OF DATA DECK:

```
G0025
P0369
G    AMD-RESEARCH
+0123
P0456
+0789
G0025
```

EXPLANATION:

```
Group 25 only
Problem No. 369 only
Group AMD - RESEARCH only
Two problem numbers
    (123 and 456)
Group 25 plus
    problem number 789.
```

3.1.3 Removing Files

There are three operations jobs which remove files from the 7600 PF directory: (1) the EXPIRE job, (2) a job to purge files with invalid account numbers, and (3) a job which executes the allocation algorithm. The EXPIRE job is run each day some time after midnight to remove files which have outlived their retention period. The expired files are purged from the PF directory without a copy being made.

After the EXPIRE job is run, a job is executed to compare the account number of each PF to the list of currently valid account numbers. Files not having a valid account are removed from the system without a copy of the file being made.

At the start of each day a job, PFALLOC is run to exercise the allocation algorithm (see section 3.1.1). This is the job that brings the PF base down to a reasonable size in accordance with the total mass storage space available on MFZ. Files removed by PFALLOC are dumped to one of three disk packs, where they are kept until the pack is needed for other dumps - usually a maximum of about three days. A list of the dumped files is kept by operations and, while files can be retrieved from these packs, this retrieval is not done routinely and is meant to be used as an emergency procedure only.

3.1.4 File Backup

A great deal of computer resources and peoples' time is spent on backing up the PF base to insure file integrity in case of failure. The whole 7600 PF base is routinely dumped to packs every night except Monday and Wednesday when 24 hours of production are scheduled. Two different sets of packs are used for

this dump, so that if a problem develops with the PF base, the old files can be reloaded from the last dump or even the previous one.

In addition to the daily dumps the packs used for production are changed every Sunday night with the old set being stored for two weeks. It is not possible to reload individual files from these production packs. The production packs can only be used to drop back the PF base to some older version, which would be done only if the daily dumps were not usable.

In addition to the daily and weekly backup efforts, there is a dump to tape of the whole MFZ PF base done twice a year. The full dump to tape fills up about 16 reels of 9-track 2400 foot tape written at 1600 bpi. These tapes are stored for some emergency use by the CSCF.

In spite of many precautions, files do sometimes get lost or destroyed and each user is ultimately responsible for being able to recreate his own file. For example, a few files might be damaged, say by an error in the disk controller, and this damage might not be detected by the system. If several weeks should elapse before the file is used, all the normal file backups will have the damaged files on them and the original good file will be lost. The older a file gets, the harder it is for most users to recreate it. Take the time to back up your files if they are costly, difficult, or impossible to recreate.

3.1.5 Accounting

PF accounting is done each night before files have been expired. Files catalogued with a retention period of 1 day or zero days (the default) are charged for one day.

3.2 6600

The PF's on the 6600's are managed as separate permanent file bases with different rules and procedures for each machine.

3.2.1 MFA

The PF base on MFA is supported with an archiving scheme and with daily and weekly backups.

3.2.1.1 Archiving

A scheme has been implemented on MFA to extend the available PF space by migrating files from disk to magnetic tape while keeping the entry for the file in the permanent file directory. Attaching a file whose image is not on disk causes a spot job to be generated to fetch the file image from tape and place it on line. Migration of files off line is taken care of by the archiving system, with files being dumped to tape only when the disks become about 80% full. The algorithm used to determine which files should be removed to disk is quite complicated. The goal of the algorithm is to minimize the number of times files must be retrieved from tape while freeing up as much disk space as possible. Thus large files which have not been used recently are the first candidates for archiving, while small, often used files remain on disk.

To insure that files backed off to tape can be read, two copies of each file are written onto separate archive tapes.

A limit on the maximum file size which will be archived has been established. Files larger than 1 million words (300 RB's) will not be archived. If the file is catalogued or extended and is over 300 RB's in size, the dayfile message to warn that the file will be purged if selected for archiving is issued:

FILE SIZE GT 300 RBS-NOT ARCHIVABLE

If the archiving algorithm selects a nonarchivable file for archiving, the file is purged. In the current scheme these large files are always chosen to be archived so that they are purged from the system each day.

The archiving scheme is meant to extend the size of the PF base but will not make the space infinite. There is an upper limit to the number of tapes that can be managed and a limit to the number of entries in the PF directory.

For BATCH jobs archiving is transparent to the job even when a file must be retrieved from tape. An ATTACH for an archived file will result in a delay in the job processing while the file is brought back on line from tape. There are dayfile messages posted to indicate that a retrieval is taking place.

An interactive ATTACH or FETCH will be delayed if the file being attached is not resident on the system. After an interactive attach of an archived file the following messages are to be expected at the user terminal:

ARCHIVE RETRIEVAL IN PROGRESS

WAITING FOR ARCHIVED FILE

WAITING FOR ACCESS TO FILE

ARCHIVE RETRIEVAL COMPLETED

The user need not stop all activity at the terminal to wait for the retrieval of the attached file. Once the message WAITING FOR ARCHIVED FILE has been received, the ATTACH or FETCH command may be aborted and the retrieval will continue. In the case of a user abort, no more messages about the archive retrieval will be sent to the terminal. To find if the archived file has been returned to disk, another FETCH or ATTACH must be issued.

If the user receives the message "HELP CALL ERROR" due to a FETCH, or "ARCHIVE RETRIEVAL ABORTED" due to an ATTACH, the user should contact the shift

supervisor to request the output from the SPOT job that loaded his file for a more complete explanation of the error.

3.2.1.2 Removing Files From the System

There are four reasons for removing PF's from the MFA PF system:

1. Invalid account number
2. Expired files
3. Files that should be archived but are too large or have NA=1 set
4. Inactivity; an archived file has not been used for more than 50 days

All purging or dumping of files is done by the archiving system job which is run at least every night and perhaps run during the day if the MFA disk storage becomes too full. The archiving system keeps a record of all the files removed from the system for the last 14 days, and the program PFGONE is provided for looking at the list of removed files. In the case of expired files an image of the file is made before it is removed from the system and the copy is saved for about one week, so the expired file may be retrieved in case of an emergency.

Files removed from the system because of inactivity are files which have been archived to tape and have not been attached for about 50 days. Files of infinite retention period with a valid account number which are less than 300 RB's (1 million words) are in no danger of being removed from the system until they have been inactive long enough to be archived and then will remain archived for 50 days.

A program PFGONE is provided to help find PF's which have been removed from the system. The program resides on the BNLU library and can be executed with the following control statements entered as a BATCH stream or entered interactively.

ATTACH(BNLU)
LIBRARY(BNLU)
PFGONE.

Input to PFGONE is a PF name or names. The program will report on whether there is a record of the specified file(s) having been removed from the system.

PFGONE will indicate if the removed file has been dumped to tape.

The input format is as follows:

ID=id, PFN=pfn, CY=nn

with one file per input record.

The ID= is required on each input record, the pfn and cycle number are optional.

If just the ID is specified all files with the indicated ID which have been removed from the system in the past 14 days will be reported on by PFGONE. For BATCH jobs the input to PFGONE should be on the INPUT file.

3.2.1.3 File Backup

Like the MF2 PF base, the MFA files are backed up on a regular schedule as part of an effort to insure the integrity of the PF base.

Each night all MFA PF's that have been created or changed during the last 24 hours are imaged to magnetic tape; this is the daily incremental dump.

In addition to the daily incremental dumps, twice a week, on Wednesday and Saturday nights, the entire PF base is imaged to tape. The full dump tapes are saved for one week; that is, there are two sets of tapes for the full dump, one used on Wednesday, and one set on Saturday.

3.2.1.4 Accounting

Accounting for the MFA PF is done each night before files have been expired. This means there is a charge for files of 1 day or 0 day retention period.

3.2.2 MFB

The PF support on MFB is aimed at providing a file base for the MFB INTERCOM user. The MFB PF's are best thought of as "resident files."

3.2.2.1 File Management

Files are removed from MFB according to the following scheme:

- i Files of more than 300 RB's are purged regardless of use.
- ii Files of between 6 and 300 RB's are purged after being inactive (not used) for 2 days.
- iii Files of 5 RB's or less will be removed after being inactive for 60 days.

The job to remove files is run each night after midnight before PF accounting is done. The job is also run when the disks become too full. The parameters 5 RB and 60 days are subject to change, with notice, depending on availability of PF disk space.

3.2.2.2 File Backup

Each night all MFB PF's created or changed during the last 24 hours are imaged to magnetic tape. This tape is retained for seven days so that in case of a catastrophic failure the files created in the past week could be restored. The full base is not dumped. This means that files over one week old, which are

on MFB only, have no backup maintained by the Center. Files can be saved on the MFA base by using the SAVEPF control statement as described in Section 5.

4. USER FILE BACKUP

The user is responsible for his own data and should not rely on the CSCF for the ultimate backup of his file. If it is time consuming or expensive to recreate a file, a copy of the file must be kept by the user himself and this copy should be stored somewhere other than the computer center. Tools are provided for dumping files to tape and reloading them.

4.1 DUMPF

The details of the DUMPF utility program are described in the SCOPE 2.1 Reference Manual, pages 6-39 to 6-50, and the SCOPE 3.4 Reference Manual, pages 5-23 to 5-36. The significant difference between the utility as described in the reference manuals and the utility as installed at BNL are:

- i Only files catalogued under the same account number as the DUMPF job will be dumped.
- ii Only nondestructive MODE 1 dumps are permitted.
- iii The DUMPF control statement requires PW=MODEL.

The Reference Manuals describe the many options provided to select files to be dumped or restored. For example, the option DA=YYDDD dumps only files catalogued, modified, or extended since a given date. Files can be dumped by permanent file name or by inactivity period. It is possible to dump only files with a given USER-ID within the permitted account. The USERQA file contains sample jobs using DUMPF.

Judicious use of these options can avoid redundant dumping/reloading of aged files, files not requiring backup, or files belonging to another user with the same account number.

Some important considerations on the use of DUMPF.

- i Files attached to a running job will not dump. If any files are attached to a job while DUMPF is running these files will not be dumped. A dump job attaches the file it is dumping so two DUMPF jobs running at the same time will cause trouble, both will be incomplete.
- ii Operations runs the DUMPF utility to do normal file management. It may happen that a user DUMPF job will abort because an operations DUMPF job is running or another user DUMPF job is running using different options.
- iii On MFA archived files are not dumped by DUMPF. A file must be disk resident before it is dumped.

The moral of this is: CHECK THE OUTPUT OF YOUR DUMPF BACKUP RUN CAREFULLY BEFORE YOU REST SECURELY. It is good practice to do an AUDIT in the same run as DUMPF, with parameters matching those of the DUMPF, so that you can compare the two outputs.

The proper application of DUMPF/LOADPF requires a certain degree of sound file management by the user. For example, if you share an account number with others and simply dump and reload without any criteria, you may reload files or cycles that another user has discarded. The use of labeled tapes as backup tapes is required; but be sure to specify a retention period (T=) on the LABEL control statement on the 7600 because the default is 999. See USERQA for examples.

4.2 LOADPF

The LOADPF routine is used to load files from a tape to which they had been dumped. The LOADPF utility is described in the same section of the refer-

ence manuals as the DUMPF utility. The one important difference in the LOADPF as implemented in the CSCF is that LOADPF loads only those files with the same account number as the LOADPF job. USERQA has sample jobs for LOADPF.

4.3 Some Good Advice and Sample Jobs

A means of establishing the integrity of a permanent file backup tape is to RETURN(DUMTAPE) (7600) or UNLOAD(DUMTAPE) (6000) and then re-STAGE or re-REQUEST as appropriate for LOADPF(CL). The CL parameter generates a complete list of all permanent files on the tape; no files will be loaded by this LOADPF job since they already exist but the tape will be read. This list should be compared with the output of the DUMPF step and an audit to insure that all files have been processed properly. An immediate reading of the backup tape may disclose tape problems (parities, missing trailer tables, etc.) before it is too late to re-run the backup. A good rule of thumb to follow is: if there is any question about the content or quality of the backup tape don't rely on it, make another. That is good defensive computing.

The following sample jobs dump all the files catalogued with account number 1234

7600

```
ANY456,STMFZ.  
ACCOUNT(urname,1234)  
STAGE(DUMTAPE,PE,POST,N,VSN=K....)  
LABEL(DUMTAPE,W,T=600)  
AUDIT(AC=1234)  
DUMPF(PW=MODE)  
RETURN(DUMTAPE)  
STAGE(DUMTAPE,PE,PRE,E,VSN=K.)  
LABEL(DUMTAPE,R)  
LOADPF(CL)
```

6600

ANY456,STMFA,NT1.
ACCOUNT(urname,1234)
REQUEST(DUMTAPE,PE,N,S,RING,VSN=K....)
AUDIT(AC=1234)
DUMPF(PW=MODEL)
UNLOAD(DUMTAPE)
REQUEST(DUMTAPE,PE,E,S,NORING,VSN=K....)
LOADPF(CL)
AUDIT(AC=1234)
7/8/9
6/7/8/9

5. THE USE OF PERMANENT FILES BETWEEN MAINFRAMES

The disks of the three mainframes are not shared, but it is possible to move files from the disks of one mainframe to the disk of another mainframe. The control statements for doing this are straightforward, but some knowledge of file type and record manager operation is required when moving files between MFZ and the 6600's. Some examples are given below. These examples and more can be found in the USERQA file on MFA.

5.1 SAVEPF, GETPF, and PURGE

The control statement SAVEPF is used to catalog a local file on a specified mainframe's PF base. The parameters for SAVEPF are the same as those for CATALOG with the addition of ST=mid where mid is the mainframe id of the system where the CATALOG will take place. When the SAVEPF statement is executed a copy of the named file is sent to the mainframe where the file is to be catalogued. The local copy of the file remains. Any change to the local file after the SAVEPF will have no effect on the file catalogued. SAVEPF, GETPF and PURGE with ST parameter are described in the SCOPE 2 Reference Manual and the SCOPE 3.4 Enhanced Station Reference Manual.

For a job to retrieve a file in the PF directory of a specified mainframe the GETPF control statement is used. The parameters of the GETPF are the same as the ATTACH with the addition of the ST=mid parameter as described above. The file is not actually shipped from the mainframe where it resides until the time the file is used (opened). When using GETPF under INTERCOM, the function which

opens the file must not be aborted. If the function is aborted while waiting on MulitMainFrame (MMF) the file will not be transferred.

A file may be purged on a linked mainframe by using the ST=mid parameter on the PURGE control statement and specifying the file (pfn,ID) along with any needed passwords. A file purged in this fashion is not attached.

5.2 Some Examples

When moving files from 7600 to 6600 the user must be aware of the different file types used by the SCOPE 2.1 and SCOPE 3.4 operating systems. It is much beyond the capacity of this section of this volume to detail the file types and their uses. Some examples of how to move PF's of common file types between mainframes are given below.

5.2.1 Card Image and Line Image Files

Files which are card image or print line image such as source statements, listings or data cards, are the most common type of files to transfer between 7600 and 6600. These files contain Z type records on the 6600. Z type records are the only type accepted as input to SCOPE 3.4 utilities such as PAGE, EDITOR, FTN, and UPDATE. On the 7600 the default files type unblocked W type records. The default tape files are I blocked.

5.2.1.1 7600 Job To Retrieve a "Card" Data File Form MFA

```
JOB1,STMFZ,TPO.  
ACCOUNT(yourname,1234)  
GETPF(DATA,RUN27,ID=yourid,ST=MFA)  
FILE(DATA,RT=Z,FL=90,BT=C)  
FTN.  
LGO.  
7/8/9
```

```

PROGRAM XX (DATA,OUTPUT,TAPE1=DATA)
READ (1,10) X,Y,Z
10 FORMAT (3F10.3)
.
.
END
6/7/8/9

```

The file card is needed to inform the record manager that the record type of file DATA is Z with all records being less than or equal to FL, 90 characters.

* * * * *

5.2.1.2 7600 Jobs To Store Line Image Files on MFA

```

JOB2,STMFZ,TPO.
ACCOUNT(your name,1234)
ATTACH(ANA,STEP1, ID=your id)
FILE(TAPE6,RT=Z,FL=140,BT=C)
ANA.
SAVEPF(TAPE6,RUNOUTPUT, ID=yourid, ST=MFA,AC=1234, RP=2)
6/7/8/9

```

The file RUNOUTPUT written by the program ANA can be attached on MFA and read by PAGE, copied to OUTPUT, read by a formatted FORTRAN read, or put into EDITOR.

* * * * *

In JOB3 the user wishes to take a deck from an OLDPL and move it to MFA so that he can edit the file.

```

JOB3,STMFZ,TPO.
ACCOUNT(your name,1234)
ATTACH(OLDPL, ID= your id)
FILE(COMPILE, RT=Z,FL=90)
UPDATE (Q)
CATALOG(COMPILE, ID=yourid, ST=MFA,AC=1234)
7/8/9
*C TRANS
6/7/8/9

```

The file COMPILE will be catalogued on MFA with a 5-day retention period. As a Z-type file, it can be used as input to the FTN compiler, PAGE command, or copied to OUTPUT (COPYSBF) to be printed, or put into EDITOR to edit the file.

* * * * *

5.2.2 Storing Data on MFA from MFZ

```
JOB4, STMFZ, TPO.  
ACCOUNT(your name, 1234)  
STAGE( IN, VSN=NB0001, PE, NR)  
COPYP( IN, HOLD)  
SAVEPF( HOLD, TESTDATA, ST=MFA, ID=yourid, AC=1234, RP=20)  
6/7/8/9
```

JOB4 is storing some test data on MFA while a user debugs his 7600 program. The file, if not too large, will stay on MFA for the indicated retention period. To use the file on MFZ, all that is needed is a GETPF, no file card. It is assumed that the tape NB0001 is default file type for MFZ, that is RT=W, BT=I.

* * * * *

5.2.3 Moving Files Between MFA and MFB

There is no problem with file types when moving files between the 6600 except for random OLDPL files.

The INTERCOM user should be aware that the GETPF does not take place until the file to be retrieved is used (opened). The file can be opened with a REWIND. The following INTERCOM commands from MFB will retrieve (DATA, myid) from MFA.

```
{COMMAND} GETPF, HERE, DATA, ID=myid, ST=MFA  
{COMMAND} REWIND(HERE)  
WAITING MMF  
{COMMAND}
```

If the REWIND is aborted with a "%A", the file will not arrive from MFA. Some users prefer to send a BATCH job to retrieve the file so that they can go on with other work while the file is transferred.

SAVPF transfers the file upon command but it too cannot be aborted if the user wants to be sure the file has indeed been transferred.

THIS PAGE
WAS INTENTIONALLY
LEFT BLANK

6. PRIVATE PACKS

The use of private packs is supported under the SCOPE 2.1 operating system on MFZ with three disk drives reserved for private pack usage. It is expected that jobs requiring private pack usage are the longer-running production type jobs. A pack mount is costly, time consuming, and causes wear on both pack and drive. Jobs requiring pack mounts have purposely been removed from consideration as express jobs.

The packs are stock items S94100 and once initialized by operations with a VSN and a set name are quite easy to use. Before a pack is accepted for use in the CSCF its use should be explained to the operations supervisor. The 844-21 single-density packs now in use hold about 10 million words. A job card parameter, YD1, is used to indicate to the system that a job requires a pack mount and is necessary in any job that will use a pack. Before any file on the pack can be accessed via an ATTACH or before a file can be assigned to the pack, a MOUNT control statement must be issued for the pack involved. The format of the MOUNT statement is

MOUNT(SN=setname,VSN=vsu)

Once the pack is mounted the addition of the SN=setname parameter to CATALOG, ATTACH, PURGE, AUDIT, DUMPF, LOADPF, and REQUEST statements will indicate to the operating system that the function requested will be carried out on the indicated pack, not on the system packs.

More detailed information on pack usage can be found in the SCOPE 2.1 Reference Manual under MOUNT and SETNAME.

The private packs are not backed up by the CSCF and should be backed up by the user using DUMPF. Before the files on the pack can be dumped, an empty (1 word) file with a pfn of DUM must be catalogued on the pack. The following job will set up the needed file. It need be run but once.

```
JOB1,STMFZ,YD1.  
ACCOUNT(yourname,1234)  
MOUNT(SN=pack set name, VSN= pack vsn)  
REQUEST,DUM,*PF,SN=pack set name)  
COPY(EOI,DUM)  
CATALOG(DUM, ID=PUBLIC,TK=DUMPF,RD=MODEL,SN=set name)  
6/7/8/9
```

Once the DUM file is on the pack with a read password of MODEL the pack can be dumped with the following job.

```
JOB1,STMFZ,YD1.  
ACCOUNT (yourname,1234)  
MOUNT(SN=pack set name, VSN=pack usn)  
STAGE(DUMTAPE,PE,POST,E,VSN=K9997/K9998/K99989)  
DUMPF(MO=1,PW=MODEL,SN=pack set name)  
AUDIT(SN=setname)  
6/7/8/9
```

Three nine track 2400 foot reels of tape written at 1600 fci will hold the contents of one full pack.