FDRAPPL

APPLICATION BACKUP

CONCEPTS & FACILITIES GUIDE

INNOVATION® DATA PROCESSING

FDRAPPL
APPLICATION BACKUP

t a b l e o f c o n t e n t s

INTRODUCTION

CONCEPTS & FACILITIES GUIDES

For more than 30 years, Innovation Data Processing has been producing high-quality Storage Management Software. Over the years, its products have evolved into today's ultra high-speed, safe, reliable storage management solutions for OS/390, z/OS, LAN and Open Systems Data.

It all started with the **FDR Storage Management Family**, of which over 5000 licenses have now been sold worldwide. The FDR Family is the complete Storage Management System for OS/390 and z/OS.

FDR has become the industry standard for fast, reliable backups of MVS OS/390 data.

ABR adds a layer of automation to the standard functions of FDR, providing advanced backup facilities like *Incremental Backup*, *Application Backup* and *Archiving*.

COMPAKTOR and **FDRREORG** further enhance the suite by adding intelligent and powerful reorganization processes, for whole DASD volumes and for Sequential, PDS and VSAM datasets.

FDREPORT provides extensive customized DASD Management Reporting to suit many needs and purposes.

FDRCLONE is an extension to ABR, providing the ability to "clone" volumes and/or datasets on a test or disaster recovery system. It includes **FDRDRP**, a utility that can reduce ABR full-volume recovery time by up to 80%.

FDRINSTANT enables FDR/ABR to take *non-disruptive backups* of offline volumes, created by the latest DASD Subsystem features like StorageTek/IBM SnapShot Copy, EMC^2 TimeFinder/BCV, HDS ShadowImage and IBM FlashCopy.

FDRPAS (FDR Plug and Swap) allows for the non-disruptive movement of OS/390 disk volumes from one disk device to another. When new disk subsystems are installed, active online disk volumes can be swapped to drives in the new subsystem without disrupting normal operations or requiring a re-IPL. This allows a 24 x 7 installation, with no window for major re-configurations and hardware changes, to install and activate new hardware.

THE FDR/UPSTREAM Family of Products builds on the strengths of the FDR Storage Management Family providing a fast, safe and reliable solution to backing up Open Systems data from file servers and workstations, across a network connection to disk or tape on the OS/390 host. If the Open Systems data is resident on an EMC^2 Symmetrix with Enterprise Storage Platform (ESP), **FDRSOS** and **FDR/UPSTREAM/SOS** products provide additional performance enhancements to the backup and restore process by utilizing high-speed mainframe channels.

IAM is Innovation's alternative to VSAM KSDS, ESDS and (as a cost option) AIX files. It eliminates VSAM performance bottlenecks and reduces VSAM file sizes by more than 50%.

FATS/FATAR and **FATSCOPY** are a set of multi-purpose tape subsystem Media Integrity tools that allow for online tape certification, verification and erasure, as well as the ability to analyze and copy tapes.

INTRODUCTION

CONCEPTS & FACILITIES GUIDES

Each of the Innovation products is described in a range of Concepts & Facilities Guides that have been created by the Innovation UK office, but which are available *free of charge* from your local office (see back cover for details).

In this particular guide, we take a look at the **FDRAPPL Application Backup** system, which is part of ABR, but which can also be licensed separately (with FDR as a prerequisite). FDRAPPL is a competing product to the ABARs component of IBM's DFSMShsm.

PART ONE provides a basic overview of the FDRAPPL Application Backup system, including some benchmarks illustrating ABR's superior performance over DFSMShsm. It also discusses briefly how the system interacts with SMS.

PART TWO provides more detailed information, with JCL examples and accompanying descriptions. FDRAPPL's highly efficient recording mechanism is also described, together with some example restore jobs.

PART THREE then illustrates the powerful reporting facilities that are available with the FDRAPPL Application Backup system, enabling users to run reports to show which datasets have been backed up by ABR.

Any comments or suggestions regarding this guide can be directed to:
support@fdrinnovation.com

PART ONE

GENERAL OVERVIEW

What is FDRAPPL?

FDRAPPL Application Backup is designed to backup all the datasets that relate to a given application, from whatever disk volumes they reside. FDRAPPL can search the MVS catalogs for datasets matching one or more name masks, quickly and easily selecting these datasets to be processed.

FDRAPPL creates one backup file per disk and then automatically stacks multiple backup files onto the output tape. Usually, a single tape can be used for output, creating one tape with all the selected data. For disaster recovery purposes, FDRAPPL can create duplicate backups, if required.

FDRAPPL uses a database to record the datasets that have been backed up. This database is called an Application Control File (ACF), and is described in PART TWO of this guide. The ACF is very compact and can record several hundred datasets in a single DASD track. A separate ACF can be used for each application and can be optionally backed up as the last file on the tape.

If datasets need to be restored from an Application Backup, the information on their location is obtained from the ACF to build a list of backup files to be dynamically allocated and read. Individual datasets can be selected for restore, or all datasets recorded in the ACF can be restored in one go, either from their most recent or an older backup.

Application recovery uses dataset (FDRDSF-type) restores for maximum efficiency and the restored datasets may be directed to a different set of volumes than those from which they were dumped, potentially spreading out over fewer (or more) volumes than they originally occupied. This is particularly useful in disaster recovery situations.

Why Use FDRAPPL?

Many application jobstreams contain a series of steps to backup application oriented datasets, usually to provide restart and recovery capabilities. These backups, which are often taken with utilities like IEBGENER or IDCAMS REPRO, can be very time-consuming and can use an unnecessary number of tapes. FDRAPPL can be used as a single-step, high-speed replacement for these backups.

Volume Backups are usually taken at a fixed time each night. For many applications requiring a specialized backup, this may not be appropriate and may not provide the recovery and restartability they need. Volume Backups also tend to have quite rigid retention rules that usually involve the expiration of backups after a given number of weeks. Again, this may not be appropriate for specialized Application Backup that may require either a short (*i.e.* days) or long (*i.e.* months/years) retention. With FDRAPPL, both the *frequency* and the *retention* of the backup are under the control of the application.

Also, FDRAPPL selects only the required datasets for an application, ensuring that the downtime for the application is reduced to a minimum. These backups are usually taken after the batch or online processing of the application has completed, providing a *timely* backup of the data.

Finally, FDRAPPL can restore applications in a fraction of the time taken by Volume Restore methods—especially when the datasets belonging to many applications are stored on the same DASD volume, as is often the case under SMS. FDRAPPL is particularly useful in disaster recovery situations, where critical applications need to be recovered ahead of less important ones which share the same DASD volume(s).

PART ONE

GENERAL OVERVIEW

Backup Integration and Performance

ABR has always been a fully integrated part of the FDR DASD Management Family. The FDRAPPL Application Backup System uses the same common 'data mover' module as the other members of the family:

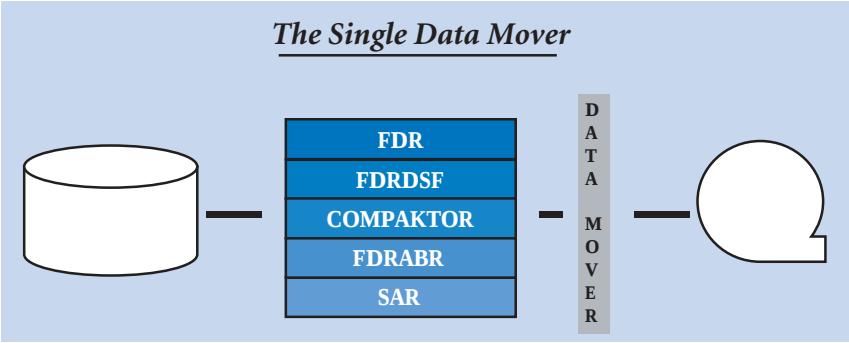

This **integration** ensures that the control statements used under ABR (including FDRAPPL) are similar to those used under the rest of the FDR DASD Management Family, providing a consistent 'feel' across the whole system. Also, as all the members of the family create the backups in the same format, each member is capable of restoring data from a backup created by another member of the family. This is an extremely useful feature, especially in disaster recovery situations.

ABR's use of the single data mover also provides benefits in backup performance. The backups taken with FDRAPPL run at the same very high speed as FDR and FDRDSF backups, and they are equally efficient in the use of critical resources such as CPU time.

An indication of the benefits of this excellent backup performance is illustrated by the following results obtained from a benchmark test comparing FDRAPPL Application Backup with the ABARs facility of IBM's DFSMShsm:

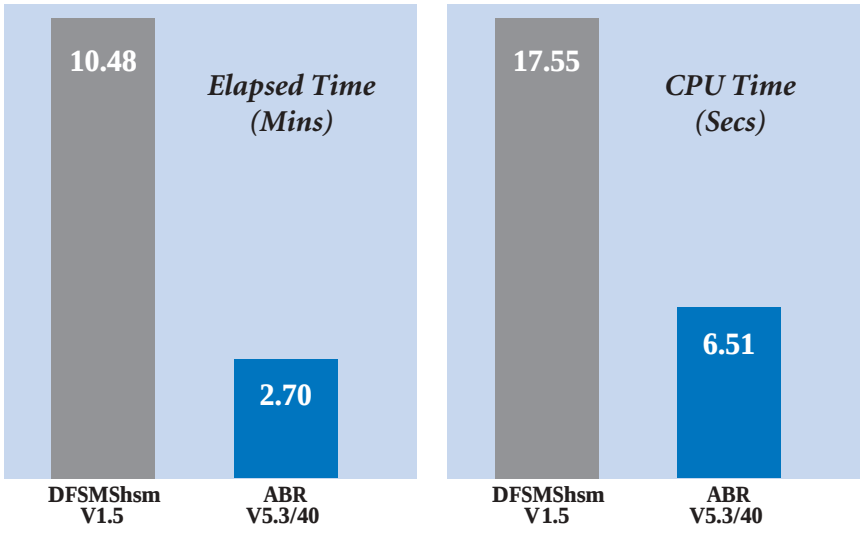

In the above test, an Application Backup was run against 1,108 non-VSAM datasets, residing on two 3390-2 DASD volumes on a RAID device. The backup was taken to a single IBM 3590-B11 tape drive on a 3590-A50 controller. To be as similar as possible to DFSMShsm ABARs, the FDRAPPL backup specified selection criteria using CATDSN (see PART TWO), forcing ABR to use the catalog to find the required datasets. Full details of this test can be obtained on request.

PART ONE

GENERAL OVERVIEW

HSDM and ExHPDM Support

Users of FDRAPPL who are also licensed for FDRINSTANT can take advantage of the **HSDM** (High Speed Data Mover) feature, which is available in the StorageTek SVA and some IBM RVA subsystems.

HSDM allows FDRAPPL to backup data in an internal compressed format. Depending on the levels of compression achieved by the SVA/RVA, the amount of data that must be read from disk and transferred to the backup media can be reduced by up to 60%. This will lead to a corresponding reduction in the *elapsed time* of the backup task (see benchmark below). When restoring data back to an SVA/RVA, FDRINSTANT also avoids the compression overhead by directly writing the compressed track images. This provides similar reductions in restore elapsed times.

FDRINSTANT also supports StorageTek's **ExHPDM** (Extended High Performance Data Mover). ExHPDM provides the ability to interleave several concurrent backups onto a *single* tape. This is not the same as writing backup files *one after the other* on the tape. ExHPDM receives the data blocks from multiple concurrent backups and writes these blocks of data to a single tape. This makes full use of the faster channel and the improved recording speed of the cartridge drive. Users who run multiple concurrent backups (*i.e.* most users!) will improve the performance of their overall backup task and potentially reduce the number of backup tapes (and drives) required to complete the process.

To demonstrate the enhanced performance that can be achieved using the combination of FDRABR and HSDM/ExHPDM, Innovation carried out a series of performance comparison tests. Full details of these tests can be obtained on request.

In the test shown below, ABR Version 5.3/40 was used to take backups of four 3390-2 disks on an RVA device, each one populated to 75% capacity with a mix of VSAM (284 datasets) and non-VSAM (1,116 datasets). The backup device was a single StorageTek 9840 tape drive.

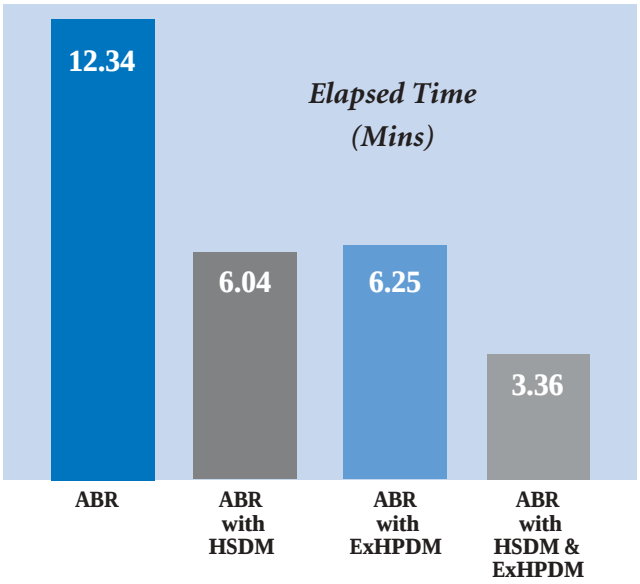

Important Note:

FDRINSTANT is required for the support of HSDM and ExHPDM within FDR and ABR, including FDRAPPL. Please see the FDRINSTANT Concepts & Facilities Guide for more details.

PART ONE

GENERAL OVERVIEW

Restore Performance

We showed earlier that the backup performance of the FDRAPPL Application Backup is equal to the other members of the FDR DASD Management Family and far superior to competing products like IBM's DFSMShsm ABARs. FDRAPPL restores are also very fast and efficient.

An indication of the performance of FDRAPPL restores can be seen in the following results that were also obtained from the benchmark test comparing FDRAPPL and DFSMShsm ABARs.

In this particular test, the same 1,108 non-VSAM datasets that were backed up in the test shown on Page 5 were restored.

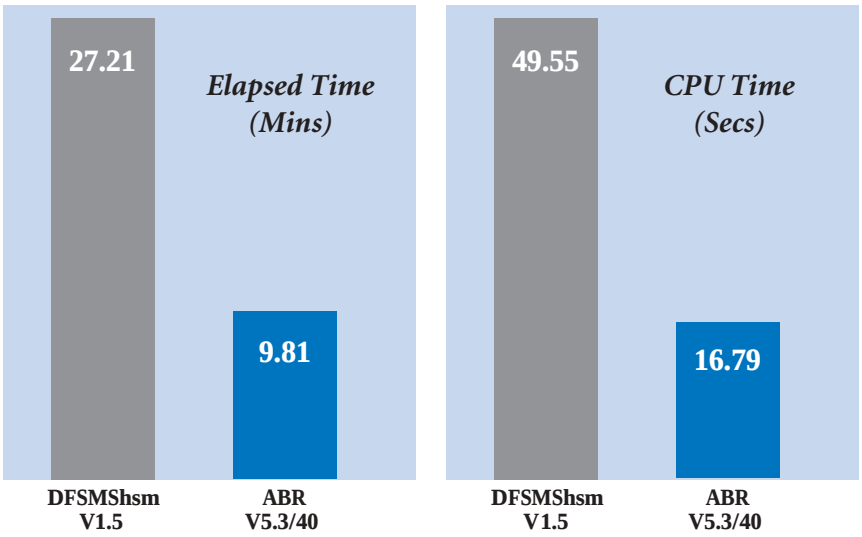

Disaster Recovery

FDRAPPL is an ideal tool to use for Disaster Recovery. It can either be used on its own, or as a complement to the FDRABR Volume Backup system, which is described in the Concepts & Facilities Guide of the same name.

For companies with a large DASD capacity, perhaps coupled to a limited recovery window, FDRAPPL can be used to secure all of the datasets belonging to one or more critical applications. These datasets can then be restored at the disaster recovery site, providing a fast and effective recovery of the application, without having to restore other non-related datasets at the same time.

By segregating the backups of individual applications, FDRAPPL allows a flexible restore approach, which can be altered depending on the time of day/week/month that the disaster occurs. For example, datasets that are used for a 'month-end' application are highly critical at month-end, but would not need to be the first ones restored if the disaster occurs at the beginning of the month. FDRAPPL allows you to restore applications in the order dictated by the current circumstances.

PART ONE

GENERAL OVERVIEW

SMS Support

MVS and OS/390 include a storage management automation function called System Managed Storage (SMS). If implemented, System Managed Storage provides centralized and automatic assignment of new datasets to pools of disk volumes, and the assignment of characteristics to those datasets that will enable them to be automatically managed.

IBM includes functions for the support of SMS in their DASD Management software DFSMSdss and DFSMShsm. However, the FDR DASD Management System also includes support for SMS, making those IBM products unnecessary.

FDRAPPL supports SMS Storage Groups as 'volume selection' criteria for Application Backup, allowing backups to be run against some or all of the online volumes in the specified SMS Storage Group. FDRAPPL can process SMS-Managed volumes and Non SMS-Managed volumes in the same job step.

During backup, FDRAPPL will record all of the SMS information (from VVR's and NVR's in the VVDS) for each dataset backed up. On a dataset restore, if the dataset already exists, the restore will be done to the pre-allocated dataset without going through the ACS and SMS volume selection routines. If the dataset is not already allocated, SMS will be invoked to decide where the dataset should be placed. ABR will present the ACS routines with the Data Class, Management Class and Storage Class values for each dataset as recorded during the backup. Alternatively those values can be nullified before going through the ACS routines using the 'NULLxxxxxxxx' parameters on the restore control statement.

There may be occasions, perhaps during disaster recovery, where there is a need to bypass this 'SMS compliance'. This can be done using certain parameters at restore time, but only by users who have the appropriate RACF authority.

Summary

FDRAPPL Application Backup provides a valuable companion to the FDRABR Volume Backup system, allowing users to take additional backups of application datasets that require special handling outside of the pre-determined frequency and retention of daily volume backups.

FDRAPPL backups use the same high-performance data mover as FDRABR Volume Backups, ensuring a fast and efficient backup and restore of application datasets. This performance is highlighted by the benchmark figures shown on pages 5 and 7.

Innovation is continually enhancing both the backup and restore functions of the system, to ensure that they not only stay abreast of the changes in DASD and Tape hardware, but that they actually *utilize* these changes to further improve performance and flexibility. The enhancements added, via FDRINSTANT, to support StorageTek's HSDM and ExHPDM (page 6) are an example of this.

FDRAPPL Application Backup continues to be the backup system of choice for Data Centers that are serious about backing up their critical applications and ensuring that they can be restored in an efficient and timely fashion.

PART TWO

BACKUP AND RESTORE EXAMPLES

Introduction

In this section of the Guide, we are going to look at some example FDRAPPL jobs that illustrate the simple JCL and control statements required to run the application backups and restores described in PART ONE. We also take a look at the Recording Mechanism used by ABR to track its backups and provide automated restore facilities.

Backup Example

The following JCL example illustrates how simple it is to run FDRAPPL to select all the datasets for a given application—in this case, from a PAYROLL application:

```
//BACKUP   EXEC PGM=FDRABR,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSPRIN1  DD  SYSOUT=*
//ARCHIVE   DD  DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1     DD  DSN=PAYBKUP.APPL1,UNIT=3490,DISP=(,KEEP),VOL=(,,,255)
//TAPE11    DD  DSN=PAYBKUP.APPL2,UNIT=3490,DISP=(,KEEP),VOL=(,,,255)
//SYSIN     DD  *

  DUMP TYPE=APPL,ARCB1DSN=PAYROLL.APPL.ACFBKP1(+1),
                 ARCB2DSN=PAYROLL.APPL.ACFBKP2(+1),
                 RETPD=14,ARCCAT=ALL
  SELECT CATDSN=PAYROLL.**
```

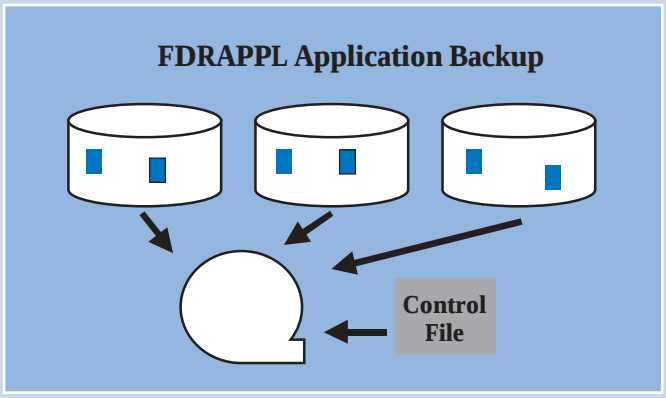

- The SELECT CATDSN statement is telling FDRAPPL to locate (via the catalogs) all datasets that begin with PAYROLL.
- The ARCHIVE DD points to the Application Control File (ACF) where the backup will be recorded (see 'The Application Control File' later). By default, a copy of the ACF will be backed up as the last file on the tape. The ARCBnDSN operands specify the names of the ACF backups. They are GDGs, allowing multiple copies of the backup to be tracked.
- The TAPE1 DD requests a 3490 to be used. The DSN is a dummy name; FDRAPPL will substitute its own name (see Page 13). Another default has been taken so that the names of the tape backup files will start with the same high-level qualifier as the ACF. The TAPE11 DD requests an additional copy of the backup, which will be sent offsite for disaster recovery.
- The two sets of backup files will be cataloged (ARCCAT=ALL) and retained for 14 days (RETPD=14). If required, a different retention could have been specified for the two separate backups.

PART TWO

BACKUP AND RESTORE EXAMPLES

Dataset Selection and Exclusion

As shown in the previous example, dataset selection in FDRAPPL can be very simple and easy-to-use. There are, in fact, several ways in which datasets can be selected (or excluded) for processing by FDRAPPL.

USING "SELECT CATDSN"

This is by far the most common method of dataset selection. The value coded on the CATDSN statement can be a fully-qualified name, *e.g.* to select a single dataset:

```
SELECT CATDSN=PAYROLL.DATA.FILE1
```

It can be a very simple mask, *e.g.* for all datasets starting with PAYROLL:

```
SELECT CATDSN=PAYROLL.**
```

Or, it can be a more complex mask to select a single dataset, or a group of datasets, by something other than just the high-level qualifier. In the following example, we are searching for the PAYROLL datasets that have a 6-character second level (where the first three characters of that level are 'SUB') and which then has any number of additional levels, ending in a final level of 'DAT':

```
SELECT CATDSN=PAYROLL.SUB%%%.**.DAT
```

Multiple CATDSN statements can be coded, if required, to build up more complex selections of datasets. For each CATDSN selection that is coded, the system catalogs are scanned for the indicated datasets. All matching datasets are then added to FDRAPPL's dataset selection list and the volumes to which they are cataloged are added to FDRAPPL's list of volumes to be processed.

USING "SELECT DSN"

'SELECT DSN' statements can be used to select *uncataloged* datasets. They can also be used to avoid the catalog search by going directly to specific disk volumes, pointed to by the 'VOL=' or 'VOLG=' operands.

The 'SELECT DSN' statement accepts the same values as CATDSN:

```
SELECT DSN=PAYROLL.DATA.FILE1,VOL=PROD01
SELECT DSN=PAYROLL.**,VOLG=PROD
SELECT DSN=PAYROLL.SUB%%%.**.DAT,VOLG=PROD
```

EXCLUDING DATASETS

'EXCLUDE DSN' statements can be used if it is necessary to exclude certain datasets from the backup that would otherwise have been selected by the CATDSN/DSN selection.

In this example, the PAYROLL datasets are being selected for backup, but it is not required to backup the library containing the PAYROLL load modules.

```
EXCLUDE DSN=PAYROLL.LOADLIB
SELECT  CATDSN=PAYROLL.**
```

PART TWO

BACKUP AND RESTORE EXAMPLES

The Application Control File (ACF)

The key to FDRAPPL is the Application Control File (ACF). When FDRAPPL is executed, details of each dataset selected for backup are recorded in the ACF.

Each record in an ACF consists of the following information:

- The name of the disk dataset that has been backed up.
- Some basic information about the dataset (*e.g.* size, type, etc).
- The name, volser(s) and fileseq of the tape file into which the dataset was backed up.
- The device type of the tape (or disk) used for the backup.
- The date the backup was taken.
- The expiration date of the backup.

Note: If two backups are created, both are recorded in the ACF in a single record.

As described in the introduction in PART ONE, FDRAPPL creates one backup file for every disk volume processed. All selected datasets on that volume go into *one* backup file on the tape.

Since many backup files may be produced in an FDRAPPL step (*e.g.* if datasets are selected from more than one disk volume), the information outlined above is required so that the restore task can find the appropriate backup file(s), without the user having to provide any of this information.

Using ACF's

There are two techniques that can be employed for the creation and use of FDRAPPL Application Control Files.

- A ***permanent*** ACF can be created for each application.
  - –The information for each new backup that is taken is appended to the data already in the ACF so that it becomes a cumulative record of Application Backups for this application.
  - –Backups are assigned an expiration date, usually calculated from a retention period (RETPD).
  - –Records for expired backups can be deleted by an ABR utility (FDRARCH).
  - –Unless otherwise instructed, a restore would select the most recent copy of each dataset.
- A ***new*** ACF can be created for each execution of FDRAPPL.
  - –Each ACF, which is usually set up as a GDG, will contain only records of datasets backed up in that execution.
  - –A restore would point to the appropriate relative generation of the ACF to obtain information about the datasets backed up in that run.
  - –Expired backups would be deleted when the ACF GDG cycled round normal GDG processing (GDG limit), thus allowing a simple specification and control of the number of backups to be retained at any one time.

PART TWO

BACKUP AND RESTORE EXAMPLES

Using a Permanent ACF

The example backup on page 9 showed the first technique, where a permanent control file was used. Here is a sample job illustrating how the FDRARCH program is used to create a permanent control file:

```
//FORMAT     EXEC PGM=FDRARCH,REGION=0M
//SYSPRINT    DD  SYSOUT=*
//ARCHIVE     DD  DSN=PAYROLL.APPL.BACKUP,UNIT=SYSDA,
//      SPACE=(TRK,(10,5),RLSE),DISP=(,CATLG)
//SYSIN       DD  *

 FORMAT  RECS=10000,USERINDEX=YES
```

The **USERINDEX=YES** operand instructs FDRAPPL to name the backup files that it creates with the same high-level index as the Application Control File, in this case 'PAYROLL'. More often than not, the ACF will be given a high-level qualifier matching that of the application datasets being backed up. This allows Application Backup files (and the ACF itself) to be easily identifiable, along with the rest of the application datasets.

The **RECS=10000** indicates the total number of ACF records that will be recorded in this ACF at any one time. This figure should be chosen based on the approximate number of datasets that will be backed up and the number of backups that will be recorded. The ACF can hold 348 records in a single 3390 DASD track, so the space required is usually quite small.

An FDRARCH utility job would periodically remove expired entries from the ACF.

Using a New (GDG) ACF

The alternative to using a permanent ACF is to create a new one for each backup. The most convenient way of controlling this is to setup the ACF as a GDG, so that each new Application Backup creates the next generation of the ACF. Here is an alternative of the example on page 9, this time using a GDG for the ACF:

```
//BACKUP    EXEC PGM=FDRABR,REGION=0M
//SYSPRINT   DD  SYSOUT=*
//SYSPRIN1   DD  SYSOUT=*
//ARCHIVE    DD  DSN=PAYROLL.APPL.BACKUP(+1),DISP=(NEW,CATLG),
//          UNIT=SYSDA,SPACE=(TRK,(10,5),RLSE)
//TAPE1      DD  DSN=PAYBKUP.APPL1,UNIT=CART,DISP=(,KEEP)
//TAPE11     DD  DSN=PAYBKUP.APPL2,UNIT=CART,DISP=(,KEEP)
//SYSIN      DD  *

  DUMP TYPE=APPL,RETPD=14
  SELECT CATDSN=PAYROLL.**
```

The ARCHIVE DD still points to the ACF. The difference is that it is now referring to a GDG (the base has already been defined) to create the next generation.

The retention period of 14 days (RETPD=14 on the DUMP statement) has been set to tie-in with the GDG limit of the control file.

Although it is not illustrated in the example, the ACF can be backed up as the last file on the tape and this backup can also be controlled as a GDG. This was illustrated in the example on page 9.

PART TWO

BACKUP AND RESTORE EXAMPLES

Backup Filenames

As described in PART ONE, the backup files created by FDRAPPL are in standard FDRDSF (dataset backup) format. Each backup file contains data from one DASD volume. If multiple DASD volumes are processed, FDRAPPL creates multiple backup files on the output tape.

Since FDRAPPL must be able to uniquely name each tape file, and must also be able to record the backup file in a way that it can easily be retrieved, it uses a special naming convention for the Application Backup files.

The name contains the disk volume serial, the date of the backup and a 'uniqueness' character in case data is archived multiple times per day, ensuring that each FDRAPPL Backup from a given disk volume will have a unique name.

The format of the FDRAPPL backup filename is:

`userindx.Vvvvvvv.bnyydddx`

where:

`userindx` If the USERINDEX=YES operand has been coded, this will be the same as the high-level index of the Application Control File being used in the backup. As described earlier, this is often the same high-level index as the actual application datasets being backed up.

`vvvvvv` Is the DASD volume serial of the disk volume from where the datasets were backed up. FDRAPPL creates one backup file for each disk volume processed in a given ABR run.

`n` Is the copy number (1 or 2) of the FDRAPPL Application Backup. FDRAPPL always creates a COPY1 backup and the COPY2 can optionally be requested by coding the additional TAPExx DD, as described on page 9.

`yyddd` Is the Julian date of the backup job (5 digits).

`b and x` There are additional qualifiers added to the backup filename to ensure that it is a unique name if multiple Application Backups are being created for the same disk volume on the same Julian date.

These qualifiers are used only if the name FDRAPPL is trying to create is already cataloged. FDRAPPL datasets do not *have* to be cataloged, although they can be, with ARCCAT=, as shown on page 9.

"b" and "x" usually have their default values of B and A.

As an illustration of the above, if a backup of our sample PAYROLL application was taken on 2000.147 and some of the datasets were backed up from volume PROD01, the resulting backup file would be called:

```
PAYROLL.VPROD01.B100147A
```

*Note: Although the above names are created by FDRAPPL, the backup JCL **must** include a DSN on the TAPEx DD statement to conform to MVS JCL conventions. As illustrated in the example job on page 9, a DUMMY dsn is coded, with FDRAPPL overriding this at the time of the backup.*

PART TWO

BACKUP AND RESTORE EXAMPLES

Restore Examples

To restore from an Application Backup, one or more SELECT statements are provided to identify which datasets are required. A special parameter, SELECT ALLDSN, can be used to select all of the datasets recorded in the Application Control File.

The JCL of the restore job will point to the Application Control File in which the backups were recorded. Since the ACF is organized in chronological order, based on the backup date, FDRAPPL will search for the requested datasets by reading it in reverse order.

If a "permanent" control file is being used, meaning that datasets have been backed up and recorded into it *multiple times*, FDRAPPL will default to selecting the most recent backup for restore. If required, earlier versions of the backup can be requested by specifying parameters such as ADATE (the date of the backup required) or OLDBACKUP (the nth oldest backup that is recorded).

If two copies of the backup were created, FDRAPPL will default to restoring from COPY=1 unless it has expired, in which case COPY=2 will be used. The second copy can be requested at any time, simply by coding COPY=2 on the restore.

If a multi-volume dataset is backed up, each volume will have a separate entry in the Application Control File, but each will be flagged to indicate it is part of a multi-volume entry. On restore, FDRAPPL will continue to search the control file until it finds all of the entries belonging to the dataset.

Restoring at the Home Site

This first example illustrates a restore at a home site, where the ACF is already on disk and available. All the files that were backed up by the example on page 9 are to be restored.

```
//RESTAPPL   EXEC PGM=FDRABR,REGION=0M
//SYSPRINT    DD  SYSOUT=*
//ARCHIVE     DD  DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SYSIN       DD  *

  RESTORE TYPE=APPL,RECAT,VRECAT,DYNTAPE
  SELECT ALLDSN
```

- The latest backup of each dataset will be restored, as recorded in the Application Control File pointed to by the ARCHIVE DD statement.
- By default, FDRAPPL will attempt to restore every dataset to its original disk volume. If this is not possible (*e.g.* the volume is full or does not exist on the restoring system), FDRAPPL can refer to a 'Restore Allocation List' for alternate volumes.
- If a dataset is SMS managed and does not currently exist on disk, SMS (or other allocation products) will influence where the dataset will be restored.
- If only certain datasets are required for restore, multiple SELECT DSNs can be coded, using the full or generic masking features highlighted on page 10.
- Operands such as ADATE or OLDBACKUP can be added to the SELECTs to request the restore of older backups.

PART TWO

BACKUP AND RESTORE EXAMPLES

Restoring for Disaster Recovery

If an Application Backup is to be used in a disaster recovery situation, it is possible that the original Application Control File, stored on disk, may no longer be available. If this is the case, a two-step procedure is required to do the restore.

STEP1 would restore the ACF from the last file on the backup tape. STEP2 would then do the restore of the selected datasets, referring to the just-restored control file. The input tape would be mounted only once, unless the backup used multiple tape volumes.

The following example job illustrates this two-step process.

```
*Restore the ACF.

//STEP1     EXEC PGM=FDRDSF,REGION=0M
//SYSPRINT   DD  SYSOUT=*
//TAPE1      DD  DSN=PAYROLL.APPL1.ACFBKP2(0),DISP=(OLD,PASS)
//SYSIN      DD  *

  RESTORE TYPE=DSF,RECAT
  SELECT ALLDSN


* Select and restore the files.

//STEP2     EXEC  PGM=FDRABR,REGION=0M
//SYSPRINT   DD   SYSOUT=*
//ARCHIVE    DD   DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//TAPE1      DD   DSN=PAYROLL.APPL1.ACFBKP2(0),DISP=(OLD,KEEP)
//SYSIN      DD   *

  RESTORE TYPE=APPL,RECAT,VRECAT,COPY=2
  SELECT ALLDSN,NVOL=DR*
```

STEP1 is using program FDRDSF to restore the most recent backup of the Application Control File, as controlled by the GDG named on the TAPE1 DD statement. This assumes that the catalog in which the GDG backup is recorded has already been restored. If it has not, the backup can still be selected ‘manually’ by coding the required DSN, VOLSER, UNIT and LABEL information on the TAPE1 DD.

The COPY2 backup of the Application Control File is being restored (ACFBKP2) as it is assumed that the COPY1 backup of the datasets and the ACF has been lost in the disaster. Since the Application Control File is the only dataset in this backup file, SELECT ALLDSN can be coded to restore it.

STEP2 is then restoring the actual PAYROLL datasets. Notice that the ARCHIVE DD statement is pointing to the control file that was restored in STEP1.

Again, the COPY2 backups are being requested (COPY=2 on the RESTORE statement) because it is assumed that the COPY1 backups are not available.

The NVOL=DR* operand has been coded on the SELECT to request that FDRAPPL restore the PAYROLL datasets to the disaster recovery DASD volumes labeled ‘DRxxxx’.

The RECAT and VRECAT parameters are coded to ensure that all the restored datasets are correctly cataloged (or recataloged) in the MVS catalogs that are active in the system where the restores are taking place.

PART TWO

BACKUP AND RESTORE EXAMPLES

Summary

As you can see from the example jobs shown in this section, the FDRAPPL Application Backup system is very easy-to-use. The backup processes can be initiated with a minimum amount of JCL and control statements, allowing a powerful and flexible selection of the datasets to be processed. FDRAPPL's highly efficient recording mechanism ensures that all backups are recorded and tracked for the appropriate amount of time, and then removed when no longer required.
The information stored in the recording mechanism provides for the automated restore processes of individual datasets and whole applications, either at the home site or for disaster recovery.

PART THREE

REPORTING FACILITIES

Introduction

In this final section of the guide, we are going to look at some of the reporting tools provided with the FDR DASD Management Family which are particularly useful to the FDRAPPL Application Backup System.

Reporting with FDRABRP and FDREPORT

The FDR DASD Management Family includes a reporting tool called **FDRABRP**, which can produce various fixed-format reports. The following example shows a report being produced against an FDRAPPL Application Control File, pointed by the ARCHIVE DD card. Details of each dataset recorded in the ACF are shown, including dataset name, DASD volume, backup date and expiry date of the COPY1 and COPY2 backups. Also shown are the volsers of the tapes containing the backups and the suffixes of the backup filenames, as described on page 13. Additional operands can be specified, if necessary, to limit the scope of the report to only certain datasets.

```
//PRINT      EXEC   PGM=FDRABRP
//SYSPRINT    DD    SYSOUT=*
//SYSOUT      DD    SYSOUT=*
//ABRMAP      DD    SYSOUT=*
//ARCHIVE     DD    DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SYSIN       DD    *

    PRINT         ARCHIVE
```

```
DATASET NAME    DISK   BKUPDT   EXPDT      ORG FM  SIZE LRECL ALLOC TY CN SUFFIX   FN VOLSER

PAYROLL.FILE1   PROD02 2000.110 2000.140   PS  VB 32760 32756    1T 0E 1  B100110A  2 TAP010
                       2000.210                                    2  B200110A  2 TAP034

PAYROLL.FILE2   PROD01 2000.110 2000.140   PS  FB 12004 12000  300T 0E 1  B100110A  1 TAP010
                       2000.210                                    2  B200110A  1 TAP034

PAYROLL.FILE3   PROD03 2000.110 2000.140   PS  VB 32760 32756 2729T 0E 1  B100110A  3 TAP010
                       2000.210                                    2  B200110A  3 TAP034

PAYROLL.FILE4   PROD03 2000.110 2000.140   PS  VB 32760 32756  111T 0E 1  B100110A  3 TAP010
                       2000.210                                    2  B200110A  3 TAP034

PAYROLL.FILE5   PROD05 2000.110 2000.140   PS  FB 12004 12000   10T 0E 1  B100110A  5 TAP010
                       2000.210                                    2  B200110A  5 TAP034

PAYROLL.FILE6   PROD04 2000.110 2000.140   PS  VB 32760 32756    5T 0E 1  B100110A  4 TAP010
                       2000.210                                    2  B200110A  4 TAP034
```

The more generalized reporting utility, **FDREPORT**, could also be used to report on datasets recorded in an Application Control File, using its DATATYPE=ARCHIVE function. Here is an example of an FDREPORT job to create the same sort of report as shown above. A more detailed description of **FDREPORT** can be found in the *'FDREPORT'* Concepts and Facilities Guide.

```
//PRINT          EXEC  PGM=FDREPORT,REGION=0M
//SYSPRINT         DD  SYSOUT=*
//SYSOUT           DD  SYSOUT=*
//ABRMAP           DD  SYSOUT=*
//ARCHIVE          DD  DSN=PAYROLL.APPL.BACKUP,DISP=SHR
//SYSIN            DD  *

    REPORT FIELD=(DEFAULTS,BKINFO)
    PRINT  DATATYPE=ARCHIVE
```

PART THREE

REPORTING FACILITIES

Reporting with SRS

The FDR DASD Management System also includes an online reporting tool called **SRS,** which stands for “Search, Report, and Services”. It allows users to search for information on datasets from a variety of sources and this information can then be displayed in an easy-to-use format. Users can optionally execute various TSO and ABR functions against the displayed datasets.

For FDRAPPL Application Backups, **SRS** can be used to display backup information, optionally showing all of the recorded backups for each dataset. In the following example, the user has provided the name of an Application Control File in the ARCDSN field and has specified ‘* *’ in the DSNAME field to request details of all datasets recorded in that ACF. Some values have been placed in the ‘Report’ column to indicate which fields are to be displayed for each dataset, and in which order.

```
------------------------- DATASET SELECTION: ABRBKUP ----- LINE 1/16 COL 4:6/10
COMMAND ===>                                                 SCROLL ===> HALF

ENTER SELECTION CRITERIA

             Read    Save    Submit    Find    Locate    Extract    Options    Help

FIELD          SELECTION VALUE                                        REPORT SORT
-----          ------------------------------------------------------ ------ ----
DSNAME    ===> **_____________________________________________________  1___  ____
VOL       ===> _____________________________________________________  2___  ____
SOURCE    ===> APPL (Catalog Volume Archive Appl Scratch Extract)      ____  ____
ARCDSN    ===> PAYROLL.APPL.BACKUP(0)
DEVTYPE   ===> _____________________________________________________  3___  ____
DSORG     ===> _____________________________________________________  4___  ____
BKINFO    ===>                                                       6___
SIZE      ===> _____________________________________________________  5___  ____
```

In response to the above request, **SRS** will search the ACF and display the requested information in a format similar to:

```
------------- DATASET LIST: DSLIST --------            LINE 1/42 COL 3:7/13
COMMAND ===>                                             SCROLL ===> HALF

 Read    Save    Find    Locate    Refresh    Next    Message    Printd    Help

COMMAND    DATASET NAME                  VOLSER DEVTYPE DSO   ALLOC   BKDATE
--------- ------------------------------ ------ ------- ---   -----   -------
          PAYROLL.MASTER1                SMS812 3390     PS        4 2000.210
restore   PAYROLL.MASTER2                SMS811 3390     PS       14 2000.210
          PAYROLL.MASTER3                SMS810 3390     PS      100 2000.210
          PAYROLL.MASTER4                SMS805 3390     PS       43 2000.210
          PAYROLL.MASTER5                SMS812 3390     PS        1 2000.210
          PAYROLL.MASTER6                SMS810 3390     PS        5 2000.210
```

To restore from a particular backup, the user simply types “RESTORE” on the line next to the backup they want restored. In the example above, the user has requested a restore of PAYROLL.MASTER2.

When the “RESTORE” command is entered, another panel is then displayed where the user can set options for the restore. They can, for example, request a rename of the dataset, or that the target volume be something other than the default (*i.e.* dataset’s the original volume).

The user then has the option to submit a batch ABR jobstream to perform the restore, or ATTACH the ABR program under TSO and execute the restore in the foreground (if the user has Tape Mount authority).

A “NOTIFY=” operand can be set so that the user’s TSO session will receive a message when the restore of each dataset has completed.

PART THREE

REPORTING FACILITIES

Summary

With any DASD Management system, it is essential that a good range of reporting facilities is provided. This reporting provides end-users with essential information about the backups of their datasets. It is also a vital tool for the DASD Storage Administrator, who needs to monitor the backup system to ensure that it is functioning correctly and that it is meeting the company's specific requirements for both home site and disaster recovery.

Corporate Headquarters

Innovation Plaza
275 Paterson Avenue, Little Falls, New Jersey 07424-1658
Tel: (973) 890-7300 Fax: (973) 890-7147
support@fdrinnovation.com *sales@fdrinnovation.com*
www.innovationdp.fdr.com

European Offices

FRANCE

191, avenue Aristide Briand
94230 Cachan

Tel: (33) 1 49 69 94 02
Fax: (33) 1 49 69 90 98
frsales@fdrinnovation.com
frsupport@fdrinnovation.com

GERMANY

Orleansstraße 4a
D-81669 München

Tel: 089-489 0210
Fax: 089-489 1355
desales@fdrinnovation.com
desupport@fdrinnovation.com

NETHERLANDS
(& Nordic Countries)

Brouwerstraat 8
1315 BP Almere

Tel: 036-534 1660
Fax: 036-533 7308
nlsales@fdrinnovation.com
nlsupport@fdrinnovation.com

UNITED KINGDOM

Clarendon House
125 Shenley Road
Borehamwood, Herts
WD6 1AG

Tel: 0208-905 1266
Fax: 0208 905 1428
uksales@fdrinnovation.com
uksupport@fdrinnovation.com

C O N C E P T S & F A C I L I T I E S G U I D E

March 2003