FDRABR

VOLUME BACKUPS

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

INNOVATION® DATA PROCESSING

FDRABR
VOLUME BACKUPS

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 **FDRABR Volume Backup** system.

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

PART TWO provides more detailed information, with JCL examples and accompanying descriptions of the *"Full Volume"* and *"Incremental Backup"* processes. ABR's highly efficient recording mechanism is also described, together with some example restore jobs, including one for the advanced *"Volume Reconstruct"* process.

PART THREE then illustrates some of the powerful reporting facilities that are provided with the FDR DASD Management Family and which are of particular interest to the FDRABR Volume Backup system. These reporting facilities enable users to see which datasets and/or volumes have been backed up by ABR.

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

PART ONE

GENERAL OVERVIEW

Introduction

There is rarely sufficient time these days to do a daily full-volume backup of every DASD volume. This is far too time-consuming and there usually isn't enough time to complete these full-volume backups within the 'window' allocated. Some users content themselves with taking full-volume backups just at the weekend, leaving themselves open to the loss of a week's worth of updates if a failure or disaster occurs late in the week.

The FDRABR Volume Backup system allows users to create *daily* backups of all of their DASD volumes (or a subset if required), but without having to backup *every* dataset *every* day. The system consists of:

- **Full-volume backups**

  These backups include all datasets residing on each DASD volume. Most users will do full-volume backups once a week, but they can be run to whatever schedule is appropriate. Innovation recommends that they be done at least once a month.

  If there is insufficient time to do all the required full-volume backups at the weekend, ABR includes a facility to select certain volumes each day for full-volume backups, doing Incremental backups (see below) on the remainder.

- **Incremental backups**

  These backups contain only the datasets that have been updated since the last FDRABR backup (either full or incremental). Also included is "descriptive data" for the volume, such as the label track, the VTOC/VTOCIX and the VVDS.

  The elapsed time of an incremental backup is usually much less than a full-volume backup (depending on the amount of data updated). This saving allows users to create daily backups of all the updated datasets on their disk volumes. Most users run incremental backups once a day, except on the day that the full-volume backup is scheduled.

With the above backups, FDRABR provides the ability to automatically:

- **Recreate entire DASD volumes** that have been lost due to hardware problems or other damage. At a disaster recovery site, the volumes that are required to run a system can be restored quickly and easily. As described in Part Two under "Volume Restore Example", the combination of full and incremental backups is sufficient to recreate an entire DASD volume exactly as it existed at the time the last backup was taken, without imposing the daily overhead of full-volume backups.
- **Restore individual datasets** by specifying only the dataset name. As described in Part Two under "Dataset Restore Example", datasets can be restored either from their most recent or from an older backup.

*Let's now take a closer look at the **integration** of the FDRABR full-volume and incremental backups, which is the key factor in ABR's stunning performance.*

PART ONE

GENERAL OVERVIEW

Backup Integration and Performance

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

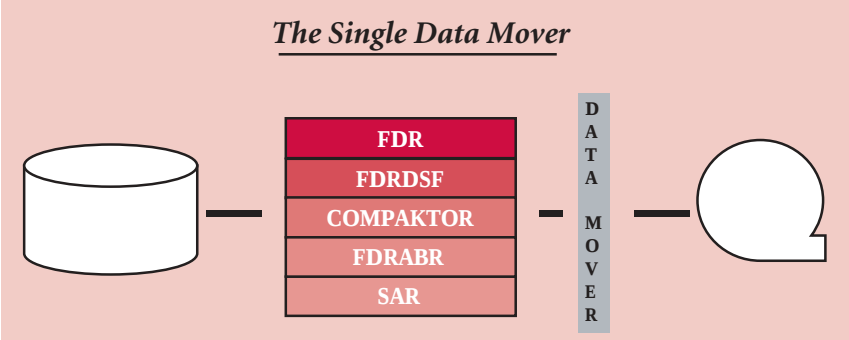

This **integration** ensures that the control statements used under ABR are similar to those used under the rest of the FDR DASD Management Family, providing a consistent 'feel' across the whole system. Also, as all the members of the family create the backups in the same format, each member is capable 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 full-volume and incremental backups 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 FDRABR and IBM's DFSMShsm:

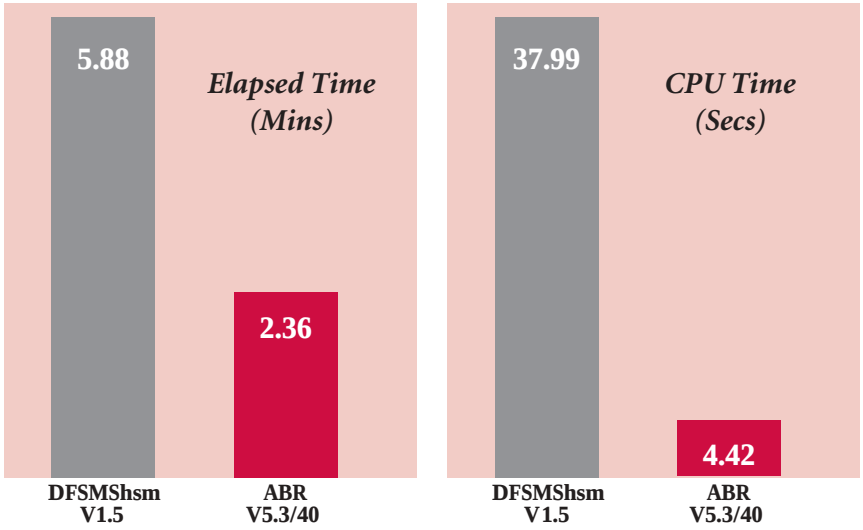

In the above test, an Incremental Backup was run against two 3390-2 DASD volumes containing 2,800 datasets, of which 304 were selected for backup (56 VSAM and 248 non-VSAM). The backups from the two disks were performed concurrently to two IBM 3590-B11 tape drives on a 3590-A50 controller. Full details of this test can be obtained on request.

PART ONE

GENERAL OVERVIEW

Utilizing FDRINSTANT

Even with the stunning performance illustrated on page 5, it may still take some hours to complete the daily backups with the standard FDRABR Volume Backup system—especially in larger shops where several hundred DASD volumes are being backed up each day.

Many companies now need to operate 24 x 7 operations and the disruption caused by daily backups, however quick, is no longer tolerable. These companies are turning to FDR Instantbackup as a solution to this problem.

FDR Instantbackup, also known as FDRINSTANT, is a separately priced member of the FDR DASD Management Family and one of its functions is to enhance the FDRABR Volume Backup system.

FDRINSTANT works in conjunction with the following:

- *TimeFinder/BCV* - EMC Symmetrix
- *SnapShot Copy* - IBM RVA (Ramac Virtual Array)
- *SnapShot Copy* - StorageTek Iceberg or SVA (Shared Virtual Array)
- *ShadowImage* - HDS 7700/7700E/9900.
- *FlashCopy* - IBM ESS 2105 (Shark)

Although their methods of implementation differ, all the above hardware features share the same common ability to create offline, duplicate copies of some or all of your DASD volumes.

FDRINSTANT then provides the ability to take “instant” point-in-time FDR, DSF and FDRABR (full-volume and incremental) backups against these *offline* duplicate volumes. This can be done while your applications (CICS, TSO, Batch etc) continue to use the real data on the original online volumes:

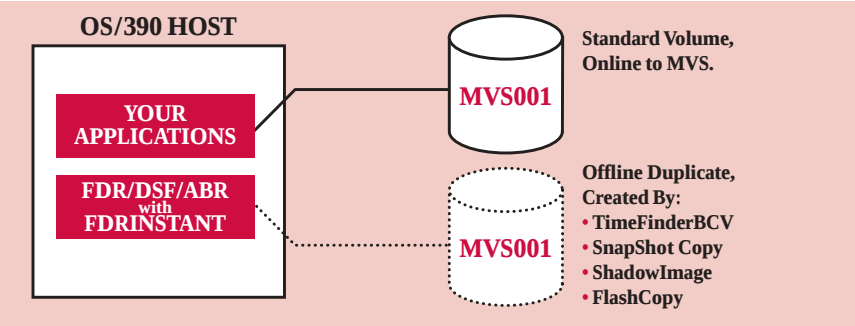

FDRINSTANT does *not* require the duplicate volumes to be brought online. It can access them across the S/390 channel while they remain *offline* to MVS. This eliminates potential problems with duplicate volsers and VTOC/VVDS and Catalog discrepancies!

Important: The FDRABR Volume Backup System has been significantly enhanced to interact with FDRINSTANT, particularly in the area of handling the MVS Update Bit during full and incremental backups. It should be noted that DFSMShsm has not been similarly enhanced and does not have an equivalent to FDRINSTANT.

PART ONE

GENERAL OVERVIEW

HSDM and ExHPDM Support

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

FDRABR 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 full volume 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.

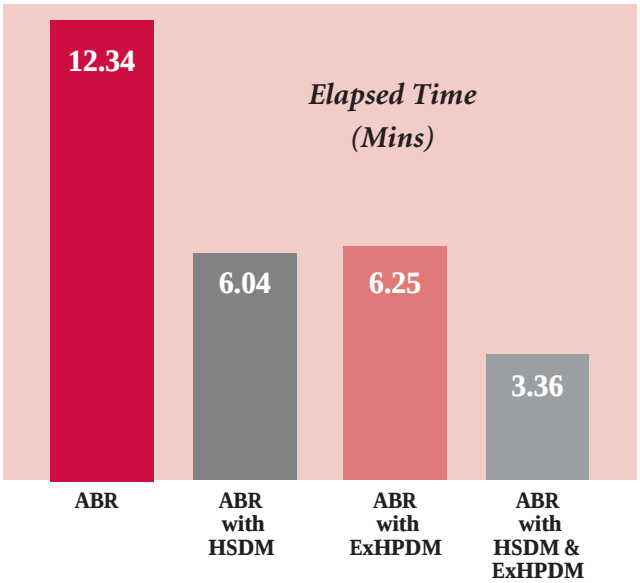

Important Note:

FDRINSTANT is required for the support of HSDM and ExHPDM within FDR and ABR, but the above savings in Elapsed Time are possible whether backing up from an offline volume, or from a standard online volume. See the FDRINSTANT Concepts & Facilities Guide for more details.

PART ONE

GENERAL OVERVIEW

Restore Performance

Part Two of this guide contains some examples of dataset and volume restores from FDRABR backups. These restore processes are also tuned for maximum efficiency.

The FDRABR volume recovery process is used by the majority of ABR users as the basis for their disaster recovery. Using the combination of the full-volume and incremental backups, they can automatically recreate whole DASD volumes at their disaster recovery site, restoring those volumes to the status they were in at the time of their most recent backup.

An indication of the performance of ABR's volume recovery process can be seen in the following results that were also obtained from the benchmark test comparing ABR and DFSMShsm. A full-volume recovery of two 3390-2 was done using a full-volume backup, together with the incremental backup described in the benchmark on page 5.

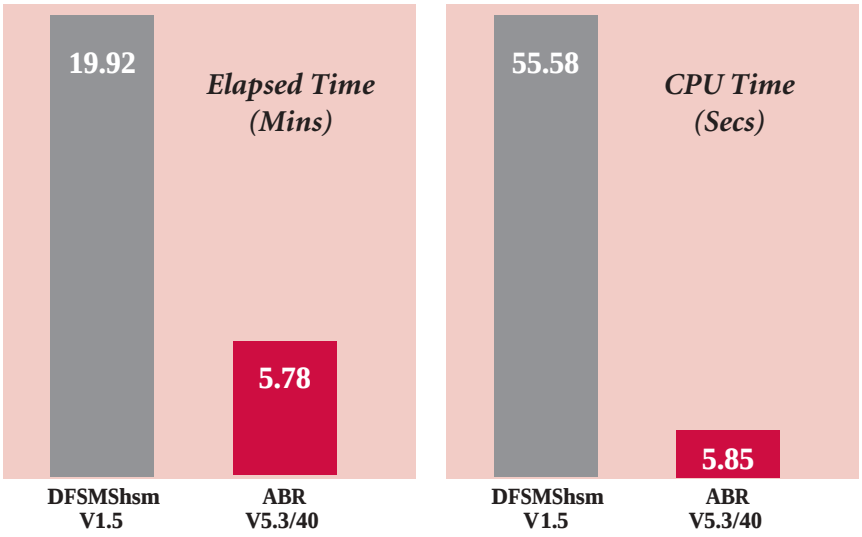

Disaster Recovery

When restoring a large number of DASD volumes—perhaps in a disaster recovery—it is important to restore as many volumes *concurrently* as the number of tape drives available. This ensures that the recovery process is completed as quickly and efficiently as possible.

ABR users typically employ *multiple* ABR Volume Recovery jobs, restoring their volumes in parallel. This method of recovery has been employed for many years by the majority of ABR customers, and it is the basis of their tried-and-tested disaster recovery plans.

However, with the advent of high performance tape drives like the IBM Magstar and StorageTek 9840's, the capacity of cartridges is continuing to increase. As new users exploit this capacity, (which ABR makes it easy to do), it is not uncommon to find the backups from multiple disks contained on the same cartridge—even the weekly full volume backups!

This consolidation of many backup files onto a single cartridge can cause serious contention problems when multiple disks are restored at the same time. The cartridge containing the multiple backups can only be mounted on one drive at a time, so other restore jobs that require the same cartridge have to wait until it becomes available.

Innovation recognizes this problem and has introduced a new member to the FDR DASD Management Family called FDRCLONE, which includes a component called FDRDRP.

PART ONE

GENERAL OVERVIEW

Utilizing FDRCLONE & FDRDRP

FDRDRP is a sub-component of FDRCLONE, which is a *cost-option* to the FDR DASD Management System. It provides an additional layer of automation to the FDRABR full-volume restore process by controlling the initiation of *concurrent* volume restores, and by passing tapes and drives between the restore tasks. It allows many disks to be restored in parallel, but ensures that input tapes are mounted a minimum number of times—often only once. This can have a dramatic effect on the overall time required to complete multiple FDRABR full-volume restores.

FDRCLONE itself also provides further options for restoring data from the FDRABR Volume Backup system. It restores datasets from the FDRABR Volume Backup system *on demand* as they are requested by online users or batch jobs. This allows users to restore only the datasets that are required at a disaster recovery site (or in a test/development system), instead of restoring *all* datasets across *all* DASD volumes, including those that will never be required.

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 which 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.

The FDRABR Volume Backup System supports the attributes of SMS Management Classes that are pertinent to the backup process. ABR also offers other selection/exclusion options, which are not available through Management Classes, but which can be used in both SMS and Non-SMS environments (*e.g.* DSORG).

ABR also supports the SMS Storage Group as a selection criterion. This allows the full-volume and incremental backups to be run against some or all of the online volumes in the specified SMS Storage Group. ABR can process SMS-managed volumes and Non SMS-managed volumes in the same job step.

During a backup of SMS controlled datasets, ABR records all of the SMS information (from VVR's and NVR's in the VVDS) for each dataset backed up. On a restore, if a dataset already exists, it will be restored to its pre-allocated space on the volume without going through the SMS volume selection routines. If it is not already allocated, SMS will be invoked to decide where it 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.

During a volume recovery, ABR will ensure that SMS controlled volumes can only be restored to another volume under SMS control. The volser of the 'input' SMS volume will be written as the volser of the output SMS volume. In disaster recovery situations, ABR will allow the reconstruction of SMS volumes to non-SMS volumes if the recovery is being run under a non-SMS system. This facility requires an appropriate level of RACF authority.

PART ONE

GENERAL OVERVIEW

Summary

The FDRABR Volume Backup system automates the backups of DASD volumes. The full-volume and incremental backups taken by the system are fast, efficient, low-resource using processes, as illustrated by the benchmark figures on page 5.

ABR can automatically recover individual or groups of datasets from these backups, and it can recover whole DASD volumes in the event of a disaster.

Innovation is continually enhancing both the backup and restore functions of the system, to ensure that it not only stays abreast of the changes in DASD and Tape hardware, but that it actually *utilizes* these changes to further improve performance and flexibility. The enhancements added by the support in FDRINSTANT for TimeFinder/BCV, SnapShot Copy, Shadowimage and FlashCopy is an example of this. Likewise, the support for HSDM and ExHPDM, and the exciting new features in FDRCLONE.

The FDRABR Volume Backup System continues to be the backup system of choice for Data Centers that are serious about backing up their critical corporate data *and* ensuring that it 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 FDRABR jobs that illustrate the simple JCL and control statements required to run the various types of 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.

Volume Backup Example

Here is an example of an FDRABR full-volume backup of a large number of DASD volumes.

ABR Full-Volume Backup

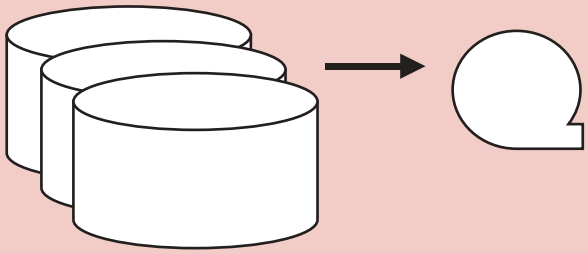

```
//BACKUP     EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT    DD    SYSOUT=*
//SYSPRIN1    DD    SYSOUT=*
//SYSPRIN2    DD    SYSOUT=*
//SYSPRIN3    DD    SYSOUT=*
//TAPE1       DD    UNIT=3590-1,DSN=FDR1,DISP=(,KEEP),VOL=(,,,255)
//TAPE2       DD    UNIT=3590-1,DSN=FDR2,DISP=(,KEEP),VOL=(,,,255)
//TAPE3       DD    UNIT=3590-1,DSN=FDR3,DISP=(,KEEP),VOL=(,,,255)
//SYSIN       DD    *

   DUMP       TYPE=FDR,ONLINE,DSNENQ=USE
   EXCLUDE    ALLDSN,VOLG=WORK
```

Some points to note about the above example job:

- The DUMP TYPE=FDR statement instructs ABR to take a full-volume dump of the selected volumes. The combination of the ONLINE parameter (on the DUMP statement) and the subsequent EXCLUDE statement instructs ABR to select all DASD volumes that are online to the MVS system, with the exception of the WORK volumes.
- Three TAPEx DD statements have been provided so that three disk volumes will be concurrently dumped. VOL=(,,,255) is specified in case any individual disk volume requires over 5 tapes. As each tape drive becomes available at the end of backing up a DASD volume, ABR will piggyback the next disk volume as the next file on that tape. The DSN's coded on the TAPEx DD's are dummy values—ABR will automatically allocate its own tape dataset names as the backups are created (see 'Recording Mechanism' later).
- The DSNENQ=USE parameter causes ABR to issue an informational message for any dataset that was OPEN at the time of the backup. These datasets will still be backed up.
- After the backup of each volume is completed and recorded, the MVS Update Bits will be turned off for all datasets on the volume.

PART TWO

BACKUP AND RESTORE EXAMPLES

Incremental Backup Example

This next example illustrates an FDRABR incremental backup of some selected DASD volumes.

ABR Incremental Backup

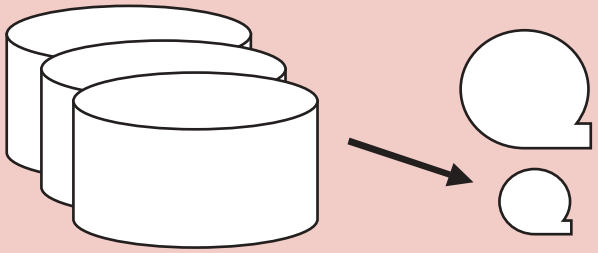

```
//BACKUP    EXEC PGM=FDRABR,REGION=0M
//SYSPRINT   DD  SYSOUT=*
//SYSPRIN1   DD  SYSOUT=*
//SYSPRIN2   DD  SYSOUT=*
//SYSUDUMP   DD  SYSOUT=*
//TAPE1      DD  UNIT=3590-1,DSN=FDR1,DISP=(,KEEP),VOL=(,,,255)
//TAPE11     DD  UNIT=3590-1,DSN=FDR11,DISP=(,KEEP),VOL=(,,,255)
//TAPE2      DD  UNIT=3590-1,DSN=FDR2,DISP=(,KEEP),VOL=(,,,255)
//TAPE22     DD  UNIT=3590-1,DSN=FDR22,DISP=(,KEEP),VOL=(,,,255)

//SYSIN      DD  *

   DUMP      TYPE=ABR
   MOUNT     VOLG=SYS
   MOUNT     VOLG=PROD
```

Some points to note about the above example job:

- The DUMP TYPE=ABR statement instructs ABR to take an incremental backup of the selected volumes.
- The MOUNT statements tell ABR to select just the SYS and PROD volumes for processing.
- Two disk volumes will be processed concurrently, creating multiple backup datasets on the tapes (one per disk volume).
- The TAPE11 and TAPE22 DD statements request that ABR creates duplicate backups, identical to those created by the TAPE1 and TAPE2 DD's. These backups would be placed offsite as a provision for disaster recovery.
- All datasets that have been updated since the last ABR full-volume or incremental backup of that DASD volume will be selected. After the backup has completed, their MVS Update Bit will be switched off.

As described in the Introduction in Part One, the usual process is to take full-volume backups at the weekend and incremental backups throughout the week, but this is not mandatory. The full-volume backups can be taken at *any* time, allowing a staggering of the process if required.

Innovation recommends that the full-volume dumps are done at least once a month.

PART TWO

BACKUP AND RESTORE EXAMPLES

The Continuing Cycle

Assuming that a standard process is in effect, and full-volume backups are being taken at the weekend, followed by incremental backups throughout the week, the backups of a sample volume (PROD01) would look something like this after one week:

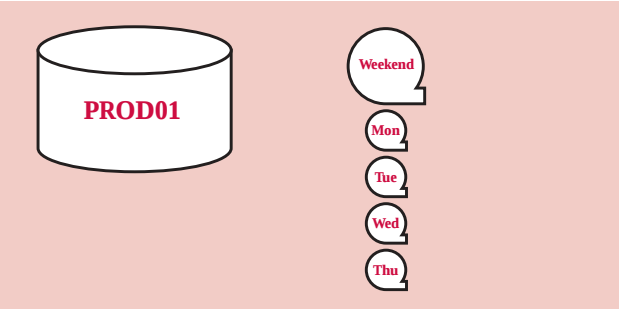

The ABR Administrator can control the number of week's worth of backups that will be retained for each DASD volume being secured by ABR. So, as additional backups are taken of our example PROD01 volume, the collection of retained backups would increase week-by-week, as shown here:

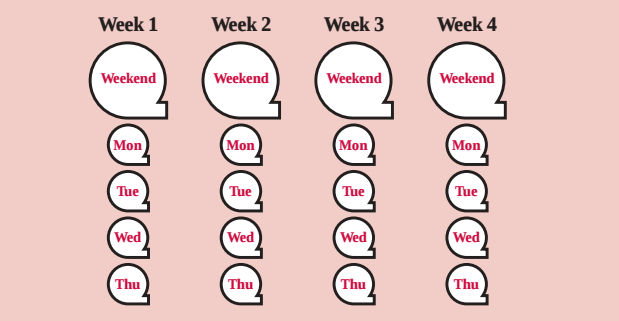

If the retention of the backups has been set to 4 weeks, when the full-volume for the fifth week is created, the full-volume and all the incremental backups for the first week will be removed automatically:

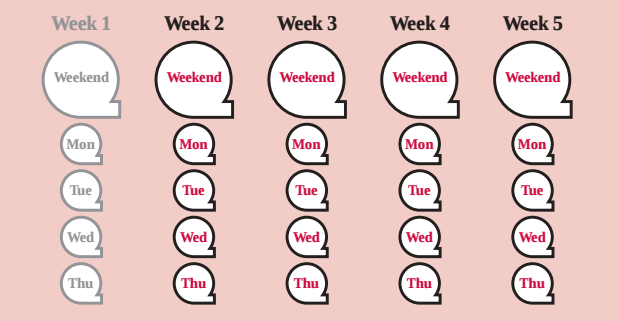

PART TWO

BACKUP AND RESTORE EXAMPLES

The ABR Recording Mechanism

The FDRABR Volume Backup System automatically records the details of all its backups and then tracks them for as long as required. Unlike other DASD Management systems, it does not rely on a complicated recording mechanism (with large control files) to do this.

Instead, it uses a simple but effective method to reduce the disk space requirements and I/O overheads. Each backup taken from a DASD volume, be it a full-volume or an incremental backup, is called a **CYCLE**.

As shown below, the full-volume backup is always CYCLE 0 and its subsequent incrementals are referred to as CYCLE 1 to nn, where nn is a maximum of 63. The full-volume backup and all of the incrementals taken up to the next full backup are known collectively as a **GENERATION**.

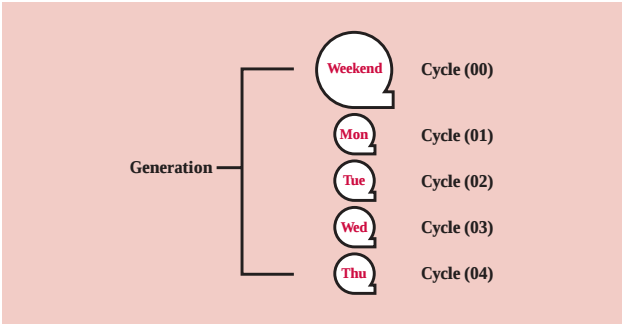

Then, as illustrated in the diagram below, each individual backup is identified by a unique GENERATION and CYCLE number. As explained earlier, the ABR Administrator can decide how many weeks (Generations) to retain within the recording mechanism for each individual disk volume.

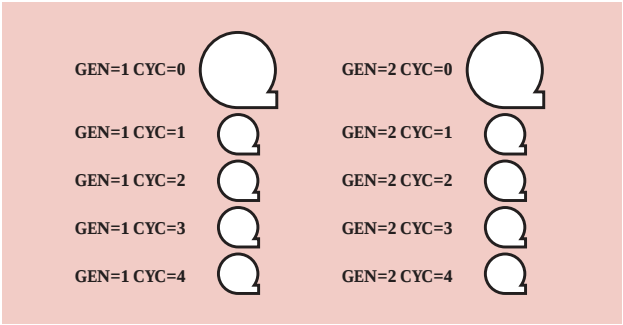

Each backup of each disk volume is then recorded in a standard ICF VSAM user catalog with the following naming convention:

- High Level Qualifier - FDRABR (*A default that can be changed*)
- Middle Qualifier - Vvvvvv (*Where vvvvvv is the DASD volser*)
- Low Qualifier - CnGGGGCC

```
|  |     |
|  |     |_____ Cycle
|  |___________ Gen
|______________ Copy
```

Thus, the COPY 1 backup created from our PROD01 volume for GEN=0003, CYCLE=04 would be named—FDRABR.VPROD01.C1000304

PART TWO

BACKUP AND RESTORE EXAMPLES

As shown in the diagram below, it is these unique backup file names that are cataloged in the ABR ICF Catalog, along with the tape volser(s), the file number and the device type. This information is used during an automated restore, so that it does not have to be provided by the user.

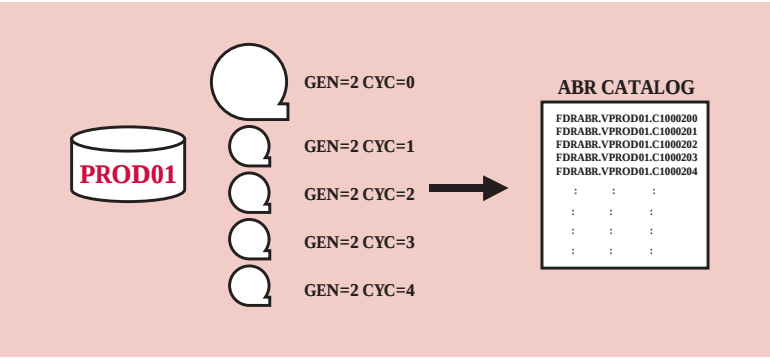

Details of the actual *datasets* that were backed up from PROD01 are not recorded in the ABR Catalog. This is because it would be a massive processing (and disk space) overhead to record in the catalog an entry for *each* backup of *every* dataset *each* day. Instead ABR uses a clever, yet simple way of recording this information.

When a volume is set up for ABR use, a 'model' DSCB is created in the VTOC on the volume. This model is just an entry in the VTOC and it does not occupy any additional space on the disk.

Within this model, ABR records information about the backup status of that volume, including the current GENERATION and CYCLE numbers.

Then, in each *dataset's* DSCB in the VTOC, ABR utilizes two unused bytes to record backup information for that particular dataset. This information includes the last CYCLE that the dataset was backed up on. It is not necessary to record the last GENERATION because all datasets are backed up at least once in the current GENERATION—the Cycle 0 full-volume backup.

Optionally up to 13 previous cycles ('Oldbackups') can also be recorded—see '*Automated Restores'* below.

Because every dataset already has a DSCB in the VTOC, and the ABR model DSCB uses no additional space on the volume, the size of ABR's control file mechanism (and the overhead of accessing it) is kept to an absolute minimum.

This is one of the critical factors in ABR's performance advantage over DFSMShsm, as illustrated in the benchmarks on pages 5 and 8.

By way of comparison, ABR's control files are usually 90% smaller than DFSMShsm's, which uses a much less efficient method for recording its backups.

With the above recording mechanism in place, ABR can then provide fully automated restore facilities for individual datasets, groups of datasets, or whole DASD volumes. Let's take a look now at some example automated restore jobs…

PART TWO

BACKUP AND RESTORE EXAMPLES

Automated Dataset Restores

If individual datasets need to be restored from the FDRABR Volume Backup System, FDRABR determines the full-volume or incremental backup that contains the most recent copy of each dataset selected for restore. Restoring the datasets requires just the dataset name to be specified; ABR locates and mounts the proper backup and then restores the dataset. During the restore, the dataset can also be renamed and/or directed to a new volume, if required.

If more than one dataset (or a multi-volume dataset) is requested for restore, ABR uses the backup information for all requested datasets to construct a list of the backup datasets that must be read. In most cases, ABR will *dynamically* allocate the backup datasets (on tape or disk) and read them. Tape backups are sorted, so that if multiple backup files on the same tape are required, they will be read in order without dismounting the tape in between.

Either the most recent, or older copies of the dataset(s) can be restored, as illustrated in the example below, restoring three datatsets from backups of MVS001:

- USER.FILE1 is to be restored from its most recent backup (the default).
- USER.FILE2 will be restored from a specifically named backup, using ABR's Generation and Cycle numbers to identify the backup.
- USER.FILE3 will also be restored from an older backup. Despite the presence of the two more recent backups of USER.FILE3, the OLDBACKUP=3 parameter tells ABR to bypass these and to restore from the third oldest backup.

Sample ABR Dataset restore:

```
//RESTORE   EXEC PGM=FDRABR,REGION=0M
//SYSPRINT   DD  SYSOUT=*
//SYSUDUMP   DD  SYSOUT=*
//SYSIN      DD  *

   RESTORE   TYPE=ABR,DYNTAPE
   SELECT    DSN=USER.FILE1
   SELECT    DSN=USER.FILE2,GEN=12,CYCLE=2
   SELECT    DSN=USER.FILE3,OLDBACKUP=3
```

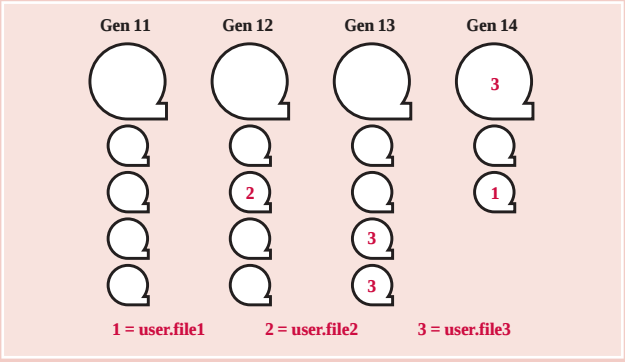

PART TWO

BACKUP AND RESTORE EXAMPLES

As you can see from the previous example, no information has been provided in the job regarding the tape volser, tape file name or tape file number where the backups reside. Also, there is no mention of the fact that the datasets reside on MVS001.

In most situations, when the most recent copy of a dataset is required (*e.g.* USER.FILE1) the user only needs to supply ABR with the name of the dataset that he or she wishes to be restored. Once supplied with the dataset name, ABR then takes the following path to do the restore:

- ABR issues a LOCATE (SVC26) for the selected dataset to find out what disk it is on. (*e.g.* PROD01)
- ABR reads the dataset's DSCB from the VTOC of that volume and obtains the LAST CYCLE value for the dataset. (*e.g.* Cycle=02)
- ABR then reads the volume's model DSCB and obtains the CURRENT GENERATION value for the volume. (*e.g.* Gen=0014)
- ABR is then able to construct the name of the tape backup file that contained that copy of the dataset. As shown on page 14, the high-level qualifier would be FDRABR, the middle qualifier would be 'VPROD01' and the low-level qualifier would be 'C1001402', producing a tape backup filename called FDRABR.VPROD01.C1001402
- ABR would then issue a LOCATE for this tape filename, which would take it to the ABR Catalog where it would obtain information such as the tape volser(s) on which the backup file resides, the tape file number and the tape device type.
- Finally, ABR would use this information obtained from the ABR catalog to dynamically allocate the tape(s) on the correct device type and do the restore of the selected dataset.

As illustrated by the other two datasets in the example job, additional control statements can be coded to request the restore from something other than the most recent version.

In the case of USER.FILE2, specific GEN and CYCLE values were coded to inform ABR to go directly to a particular backup. Coding GEN and CYCLE values relies on a knowledge of which days the dataset was backed up on. This information can be obtained by running FDRABRP or using the SRS panels, as shown later in Part Three.

An alternative, as illustrated for the restore of USER.FILE3, is the optional 'OLDBACKUP' feature. This allows previous versions of backups to be referenced by a '-nn' value (up to 13), instead of using GEN and CYCLE values.

Restoring Deleted Datasets

When a dataset is deleted, its VTOC entry (containing the ABR tracking information) would ordinarily be deleted along with it. The process described above would not then be able to obtain the information required for an automated restore.

ABR therefore intercepts the deletions of datasets (SVC 29) to check to see if they have a current ABR backup. If they do, ABR reads the backup information from the DSCB prior to the dataset being deleted and creates a special entry in the ABR Catalog, consisting of the dataset name prefixed by a '#.' (eg USER.FILE1 would be cataloged as #.USER.FILE1). The backup information for the dataset is then recorded with the new catalog entry.

Any restore request that fails the initial LOCATE for the 'real' name will be re-tried with the '#.' name. Once the last backup has expired for a deleted dataset, its special '#.' entry is removed by an ABR housekeeping job. This function also works when datasets are *renamed*.

PART TWO

BACKUP AND RESTORE EXAMPLES

Volume Restore Example

If disk failures occur, or if a DASD volume has to be restored at a disaster site, FDRABR automates the volume restore process by locating the most recent set of full-volume and incremental backups.

ABR starts off the restore process by reading the most *recently-created* incremental, restoring the label track, VTOC, VTOCIX, VVDS and any datasets on that backup. It then reads back through the preceding incrementals and the full-volume backup, restoring the most recent copy of each dataset—*once only*.

The result of this restore process is a volume that looks **exactly** like the original volume at the time of the last backup, whether it was a full-volume or an incremental backup. All the datasets are in their original locations, with the exact same allocation characteristics.

The following example job shows the JCL and control statements required to recover our sample PROD01 volume from the full-volume and incremental backups. Notice how only the volser of PROD01 has been specified. The DYNTAPE parameter tells ABR to identify from its control files the backup tapes that are required to complete the restore.

The DISK1 DD points to the output volume that will be used for the restore (SPARE1) and the CPYVOLID=YES tells ABR to rename SPARE1 to PROD01 after the contents of PROD01 have been restored to it.

Sample ABR Full-volume restore:

```
//RESTFULL  EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT   DD    SYSOUT=*
//SYSUDUMP   DD    SYSOUT=*
//DISK1      DD    UNIT=3390,VOL=SER=SPARE1,DISP=OLD
//SYSIN      DD    *

  RESTORE   TYPE=FDR,CPYVOLID=YES,DYNTAPE
  SELECT    VOL=PROD01,NVOL=SPARE1
```

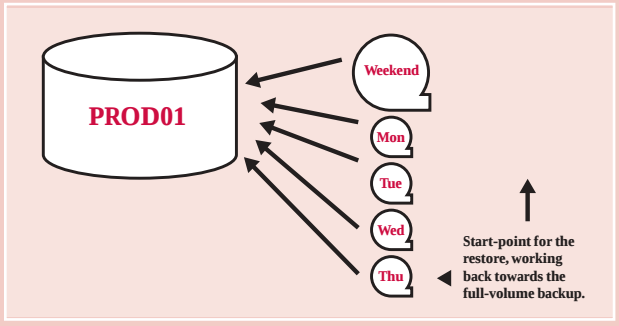

Note: This is the process that was used for the Volume Restore benchmark on page 8.

PART TWO

BACKUP AND RESTORE EXAMPLES

Summary

As you can see from the example jobs shown in this section, the FDRABR Volume Backup system is very easy-to-use. The full-volume and incremental backup processes can be initiated with a minimum amount of JCL and control statements, allowing a powerful and flexible selection of the DASD volumes to be processed.

FDRABR's highly efficient recording mechanism ensures that all backups are recorded and tracked for the appropriate amount of time, and then automatically removed when no longer required.

The information stored in the recording mechanism provides for the automated restore processes of individual datasets and whole volumes, either at the home site or for disaster recovery.

PART THREE

REPORTING FACILITIES

Introduction

Several reporting tools are provided with the FDR DASD Management System and two of them are particularly useful to the FDRABR Volume Backup System. The first of these reporting tools is called FDRQUERY.

Reporting with FDRQUERY

FDRQUERY is a statistics query program that can highlight DASD volumes that would gain the most benefit from incremental backup. It scans the volumes and estimates the number of tracks that would be dumped by a full-volume dump program (*e.g.* FDR) and compares it to the number of tracks that would be dumped on an incremental backup (*e.g.* ABR). The saving, printed as a number tracks and as a percentage, directly relates to the saving that would be made in the elapsed time and tape usage when compared to a full-volume backup.

The following example illustrates how simple it is to run FDRQUERY. A report will be produced for all online volumes beginning with PROD. A sample of the report is also shown.

```
//QUERY        EXEC   PGM=FDRQUERY
//SYSPRINT      DD    SYSOUT=A
//SYSIN         DD    *

  REPORT       BACKUP,VOL=PROD
```

FDRQUERY INCREMENTAL BACKUP SAVINGS

			ALLOCATED			INCREMENTAL			SAVINGS		
VOLSER	**DEVTYPE**	**TRACKS**	**TRACKS**	**DSNS**	**PER**	**TRACKS**	**DSNS**	**PER**	**TRACKS**	**DSNS**	**PER**
PROD01	3380	13275	7359	124	56%	3947	91	30%	3412	33	46%
PROD02	3380	13275	11223	97	85%	1634	18	12%	9586	79	86%

SUMMARY LEVEL BY DEVICE TYPE

		TOTAL	ALLOCATED			INCREMENTAL			SAVINGS		
VOLSER	DEVTYPE	TRACKS	TRACKS	DSNS	PER	TRACKS	DSNS	PER	TRACKS	DSNS	PER
TOTAL	3380	26550	18582	221	70%	5581	109	21%	13001	112	70%

PROD01 contains 124 datasets, which occupy 7359 tracks. However, only 91 of those datasets (occupying 3947 tracks) currently have the MVS Update Indicator switched on. An incremental backup against PROD01 would make a saving of 46% in the amount of tracks that would need to be backed up, when compared to a full-volume backup. This saving would have a direct effect on the elapsed time of the backup and the amount of tape required to hold the backup, reducing both by a similar percentage.

PROD02 shows even greater potential for savings. It contains 97 datasets, which occupy 11223 tracks. However, only 18 of those datasets (occupying 1634 tracks) currently have the MVS Update Indicator switched on, yielding an 86% saving on the amount of data requiring a backup.

In the final summary report, we can see that the overall savings for the two PROD volumes would be 70% when compared to doing full-volume backups. These two volumes would, therefore, show a significant savings if processed by incremental backups.

PART THREE

REPORTING FACILITIES

Reporting with FDRABRP

The FDRQUERY reports shown above are just a brief example of the extensive reporting capabilities that are provided with the FDR DASD Management Family. Another reporting tool, called FDRABRP, can be used to produce reports on the backup status of individual datsets and DASD volumes that have been secured by the FDRABR Volume Backup system.

In this first example, FDRABRP has been used to list the backups for a dataset called TEST.DSN on volume PROD01. As you can see from the resulting report, several backups of the dataset are currently being recorded by ABR. The report shows the date of each backup, as well as information about the suffix of the backup dataset name (containing the GEN and CYCLE numbers) and the tape volser(s) where each backup is stored.

```
//QUERY      EXEC   PGM=FDRABRP
//SYSPRINT    DD    SYSOUT=A
//ABRMAP      DD    SYSOUT=A
//SYSIN       DD    *

  PRINT      BACKUP,DSN=TEST.DSN
```

FDRABR DATASET BACKUP REPORT

DSN	VOLUME SERIAL	VL SQ	D/S ORG	TRK ALLOC	TRK FREE	BACKUP DATE	BK NO	TAPE SUFFIX	TAPE FILE	TAPE VOLUMES
TEST.DSN	PROD01	01	PO	12	1	2000.154	00	C1039504	0003	B90064
						2000.151	01	C1039501	0143	B90062
						2000.150	02	C1039500	0134	B90060
						2000.148	03	C1039404	0005	B90058
						2000.146	04	C1039402	0150	B90028
						2000.144	05	C1039400	0126	B90028

In the next example, FDRABRP is being used to report on information about the daily backups taken from the PROD01 volume for the current GENERATION (Gen 395).

As you can see from the resulting report, information is displayed about each backup, including an indication of whether it was a full-volume (FDR) or incremental (DSF) backup, as well as the date of the backup, the complete backup dataset name and the tape volser(s) used. If required, information about *all* generations currently being tracked can also be produced.

```
//QUERY      EXEC   PGM=FDRABRP
//SYSPRINT    DD    SYSOUT=A
//ABRMAP      DD    SYSOUT=A
//SYSIN       DD    *

  PRINT      CATLG,VOL=PROD01
```

FDRABR VOLUME BACKUP REPORT

VOLSER	GEN	CYCLE	TYPE	DUMPDATE	TAPE FILE DSN	COPY	FILE	TAPE VOLUME(S)
PROD01	395	00	FDR	2000.150	FDRABR.VPROD01.C1039500	1	1	B80223,B80247
		01	DSF	2000.151	FDRABR.VPROD01.C1039501	1	1	B80269
		02	DSF	2000.152	FDRABR.VPROD01.C1039502	1	2	B80269
		03	DSF	2000.153	FDRABR.VPROD01.C1039503	1	13	B80273
		04	DSF	2000.154	FDRABR.VPROD01.C1039504	1	29	B80278
		05	DSF	2000.155	FDRABR.VPROD01.C1039505	1	34	B80283

PART THREE

REPORTING FACILITIES

Reporting with ISPF, SRS and FDREPORT

The FDR DASD Management System also includes a full set of **ISPF** panels. The primary option menu for these panels is shown below. Among other things, these panels support the invocation of the FDRABRP reports shown above (option 1) and the restore of datasets from FDRABR Volume Backups (option 2).

Users can use the **ISPF** panels to submit their own ABR jobs to restore specific datasets, or they can add requests to a "Remote Queue" dataset for later processing. The **ISPF** dialogs also support adding *backup* requests to a Remote Queue dataset (option 4).

```
---------- FDR TOTAL DASD MANAGEMENT SYSTEM -- FDR PRIMARY OPTIONS MENU -------
OPTION  ===>
   1  REPORTS      - ABR REPORTING FUNCTIONS
   2  RESTORE      - ABR dataset RESTORE
   3  ARCHIVE      - ABR dataset ARCHIVE OR SUPERSCRATCH
   4  BACKUP       - ABR dataset BACKUP
   5  REMOTE Q     - ABR REMOTE QUEUE UTILITY FUNCTIONS
   C  COMPAKTOR    - COMPAKTOR MAP AND SIMULATION REPORTS
   R  RELEASE      - COMPAKTOR RELEASE
   I  INSTALL      - INSTALLATION AND MAINTENANCE OF FDR AND OPTIONAL PRODUCTS
   J  JCL PARMS    - SPECIFY FDR JCL AND SYSOUT DEFAULTS FOR SUBMITTED JOBS
   K  FORMAT       - MODIFY FORMAT OF GENERATED REPORTS
   M  MESSAGES     - FDR MESSAGES AND CODES QUERY FACILITY
   Q  QUERY        - FDR/ABR STATISTICS QUERY
   S  SRS          - SEARCH, REPORT, SERVICES DIALOG
   T  FDRTSEL      - BACKUP FILE MANAGEMENT UTILITY
```

Although the ISPF RESTORE panels are quite useful, users may prefer to use the **SRS** dialog (option S above). **SRS** stands for "Search, Report, and Services" and it allows users to search for information on datasets from a variety of sources. Information about the datasets can then be displayed in an easy-to-use format. Users can optionally execute various TSO and ABR functions against the displayed datasets.

For FDRABR Volume Backups, **SRS** can be used to display backup information for datasets, optionally showing all of the recorded backups for each dataset, including the date of the backup. Then, by simply typing a RESTORE command next to a displayed dataset, users can request that a dataset be restored from a particular backup. Users can also type a BACKUP command next to the dataset to request that it be included in the next incremental backup of the volume on which it resides.

In the following example, the user has filled in several dataset name masks on the main **SRS** selection panel and has specified several attributes of the selected datasets to display, (including ABR backup information). The user has requested that all recorded backups of each dataset be displayed (OLDBACKUP=ALL).

```
------------------------- 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    ===> bab.*.jcl  bab.ac**_____________________________________ 1___  ____
VOL       ===> ________________________________________________________ 2___  ____
SOURCE    ===> CATALOG (Catalog Volume Archive Appl Scratch Extract)   ____  ____
CATALOGN  ===> ________________________________________________

DEVTYPE   ===> ________________________________________________________ 3___  ____
DSORG     ===> ________________________________________________________ 4___  ____
BKGEN     ===> ________________________________________________________ 6___  ____
BKCYCLE   ===> ________________________________________________________ 7___  ____
BKINFO    ===>                                                          8___
OLDBKUP   ===> ALL____________________________________________________ ____  ____
SIZE      ===> ________________________________________________________ 5___  ____
```

PART THREE

REPORTING FACILITIES

In response to this request, **SRS** will search the system catalogs for the datasets requested and display their backup information in a format similar to:

```
---------------------------- DATASET LIST: DSLIST -------- LINE 1/20 COL 3:9/12
COMMAND ===>                                                 SCROLL ===> HALF

3 datasets SELECTED.

         Read    Save    Find    Locate    Refresh    Next    Message    Printd
Help

COMMAND    DATA SET NAME               VOLSER DEVTYPE ALLOC DSO BKGN BKC BKDATE
---------  --------------------------- ------ ------- ----- --- ---- --- ------

           BAB.JCL.CNTL                TSOWK0 3390     150 PO   230    1 2000.162
                                                              230    0 2000.159
                                                              229    4 2000.158
                                                              229    2 2000.156
                                                              229    1 2000.155
                                                              229    0 2000.152
Restore                                                       228    3 2000.151
                                                              228    2 2000.150
                                                              228    1 2000.149
                                                              228    0 2000.145
                                                              227    3 2000.143
                                                              227    2 2000.142
                                                              227    0 2000.138
                                                              226    4 2000.137

Restore    BAB.AC.DATA                 IDPLB3 3390       1 PS   721    0 2000.159
```

Some points to note about the above:

- For the dataset called BAB.JCL.CNTL, the most recent 13 backups are displayed, with the generation (BKGN), cycle (BKC) and backup date (BKDATE).
- 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 BAB.JCL.CNTL from BKGN=228, BKC=3.
- For the dataset called BAB.AC.DATA only one backup is shown. The user has typed “RESTORE” next to it to restore this most recent backup.
- 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 during the restore. Or they can request that the target volume be something other than the default (i.e. the dataset’s 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).
- Alternatively, the user can add the restore request to the ABR Restore Remote Queue.
- A “NOTIFY=” operand can also be set so that the user’s TSO session will receive a message when the restore of each dataset has completed.

The FDR DASD Management System also includes a general free-format reporting tool called **FDREPORT** which, among other things, can report on backups taken by the FDRABR Volume Backup system. For a fuller description of **FDREPORT** (and further information on SRS), please refer to the *‘FDREPORT’* Concepts and Facilities Guide.

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.

INNOVATION®
DATA PROCESSING

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

UK

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

February 2003