FDR STORAGE MANAGEMENT FAMILY

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

FDR
STORAGE MANAGEMENT FAMILY

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

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 are 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).

Several of the Guides describe in detail some of the key features of the FDR Storage Management Family (*e.g.* Volume Incremental Backup). However, in this particular guide, we are going to take an overall look at *all* of the components and modules of the FDR Storage Management Family.

(Note: Some of the modules described in this Guide are available as separate cost-options. Please contact Innovation for further information).

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

PART ONE

GENERAL OVERVIEW

Storage Management for Today's OS/390 and z/OS Systems

The FDR Storage Management System provides a high-performance, reliable and cost-effective set of tools, designed to meet the demands of today's increasingly diverse OS/390 and z/OS storage. The various modules that are available within the system can either be purchased separately or collectively, allowing Storage Managers to selectively choose and implement the facilities required to meet their own specific needs.

The FDR Storage Management System includes several options for **Backup & Restore**:

- *Volume Incremental Backup*—Through a combination of daily incremental backups, interspersed with regular full volume dumps, you can take frequent backups of *all* your DASD volumes (or a subset, if required), without having to backup *every* dataset *every* day. From these backups, individual datasets or groups of datasets can be restored (simply by being named), or entire DASD volumes can be automatically recreated.
- *Application Backup*—If required, you can also backup groups of related application datasets from whatever disk volume(s) they reside. The catalogs are searched to locate the datasets that match one or more name masks. Information about the datasets is recorded in a control file, which can be optionally copied onto the end of the backup tape. If a restore is required, information within the control file is used to provide an automated and flexible restore of some or all datasets, either from the most recent or an older copy, and either to the same name or a different one.
- *Standard Volume/Dataset Backup*—In addition to the fully automated Incremental and Application backups described above, it is also possible to take standard (*i.e.* non-managed) full-volume or dataset backups, perhaps to meet special, ad-hoc requirements.

All of these backup options are described in more detail in PART TWO of this Guide, which also provides benchmark test results highlighting the FDR Family's performance advantage over IBM's DFSMShsm.

In addition to the raw performance, most of the above backup facilities can also take advantage of the optional FDRINSTANT component, which facilitates the creation of **Point-in-Time, Non-Disruptive Backups**, utilizing the offline duplicate volumes created by DASD Subsystem features like SnapShot Copy, TimeFinder, FlashCopy and ShadowImage, as well as Innovation's own FDRPAS (see "Non-Disruptive DASD Volume Movement" on page 6).

The StorageTek features of ExHPDM and HSDM are also supported during most backups and can greatly improve the performance and efficiency of the backup tasks.

PART THREE of this Guide describes how some or all of the above backups can be used as part of a comprehensive **Disaster Recovery** plan. We also take a look at the FDRDRP component, which is a more recent addition to the FDR Storage Management family, designed to simplify and speed up the recovery process and reduce tape drive contention.

Of course, Storage Management is about more than just Backup/Restore and Disaster Recovery, so the FDR Storage Management System also provides some essential tools to help Storage Managers get the very best out of their DASD.

PART ONE

GENERAL OVERVIEW

The following facilities are described in more detail in PART FOUR of this guide:

- *Data Migration & Recall*—The **Data Migration** system provides an excellent way to control the consumption of DASD by ensuring that expired or unwanted datasets are either deleted or moved to a less expensive medium (*e.g.* tape). DFSMS Management Class information can optionally be used to control the selection of datasets to be migrated. An auto-recall facility is provided to ensure that migrated datasets, which are subsequently referenced by batch jobs or online users, are recalled quickly and efficiently, and without causing unwanted delay or disruption.
- *Dataset Copy/Move*—The FDRCOPY component of the family can quickly and efficiently **Copy or Move** datasets from one DASD volume to another, consistently outperforming utilities like IEBCOPY.
- *Dataset Reorganization*—The FDRREORG component provides highly efficient **Reorganization** services for **VSAM**, **IAM** and **PDS** datasets, outperforming utilities like IDCAMS and IEBCOPY. Intelligent selection criteria ensure that only the files that need reorganization are processed.
- *DASD Volume Reorganization*—Even on today's RAID devices, multiple-extent datasets and fragmented freespace can still lead to Sx37 abends and other related problems. The Compaktor component of the FDR Family provides an intelligent and highly efficient way of dealing with these issues via an efficient and powerful **DASD Volume Reorganization**. Typically, the FASTCPK option can consolidate the freespace and merge multi-extent datasets (including VSAM and DB2) on a 3390-3 in less than 2 minutes, outperforming utilities like DFSMSdss DEFRAG. When used in conjunction with FDRINSTANT, together with SnapShot Copy on StorageTek SVA/V960's and IBM RVA's, FASTCPK can reorganize a 3390-3 in just *10 seconds!*
- *Non-Disruptive DASD Volume Movement*—The FDRPAS (FDR Plug and Swap) component 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 available for major re-configurations and hardware changes) to install and activate new hardware.

And, finally, no Storage Management System would be complete without a comprehensive set of **Reporting** tools. The following report generators are provided with the FDR Storage Management System and are described in more detail in PART FIVE of this guide:

- *Backup & Migration Simulation:* FDRQUERY allows you to simulate the effects of running Volume Incremental Backups and/or Data Migration. This allows you to assess the potential savings that can be made by using these facilities, without first having to set-up and implement them.
- *Internal Reporting:* FDRABRP can be used to gather and report on information stored within the various control files and mechanisms employed by the FDR and ABR backup facilities. With a few simple control statements, fixed-format reports can quickly be generated to show which datasets and/or volumes have been backed up, answering key questions like *when*? *where*? and *how*?
- *Storage Management Reporting:* A wider range of more general reports can be created with the FDREPORT component—a powerful free-form report generator which can gather information from a variety of sources, including VTOCs, VVDSs, Catalogs, ABR Control Files and even DFSMShsm control files.

Let's begin our tour of the FDR Storage Management System by first taking a closer look at the facilities provided for Backup and Restore…

PART TWO

DATA BACKUP

Introduction

As described in PART ONE, the FDR Storage Management Family provides various options for backing up data, catering to the increasingly diverse needs of today's OS/390 and z/OS Storage Administrators. In this section, we are going to take a closer look at these options, describing the special backup services that are available.

Sample JCL is provided to illustrate the ease-of-use of each backup option, and some benchmark figures are also included to show how the performance of each option compares to IBM's DFSMSdss and DFSMShsm.

Volume Incremental Backup

Without doubt, the most commonly used backup option within the FDR Storage Management Family is the *Volume Incremental Backup* system, which is designed to secure the contents of whole DASD volumes on a regular basis *without* the overhead of taking daily Full Volume backups.

Although the system is volume-based, individual or groups of datasets can be restored as easily as complete DASD volumes. This flexibility during restore makes the Volume Incremental Backup system an essential tool for both on-site backup *and* Disaster Recovery.

The Volume Incremental backup system is described in detail in the *"Volume Backup"* Concepts and Facilities Guide. However, here is a quick overview of how it works.

As shown is the diagram, a full volume backup (usually taken at the weekend) is supplemented by regular incremental backups, which contain copies of only the datasets that have been updated since the previous backup.

A copy of the volume's VTOC is also taken each day, together with the VTOCIX, the VVDS and any Catalogs which reside on the volume. Details of each backup (*e.g.* tape volser, filenum, fileseq, etc.) are recorded in a control file.

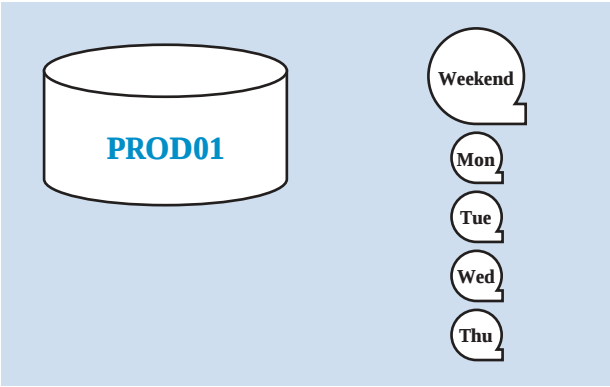

PART TWO

DATA BACKUP

Each disk that is secured by the Volume Incremental system can have several weeks of backups. The actual number of weeks to be recorded for each disk is controlled by the administrator, allowing (for example) the Production System volumes to have more week's backups than Development volumes.

As newer backups are created, the older ones are automatically removed from the recording mechanism.

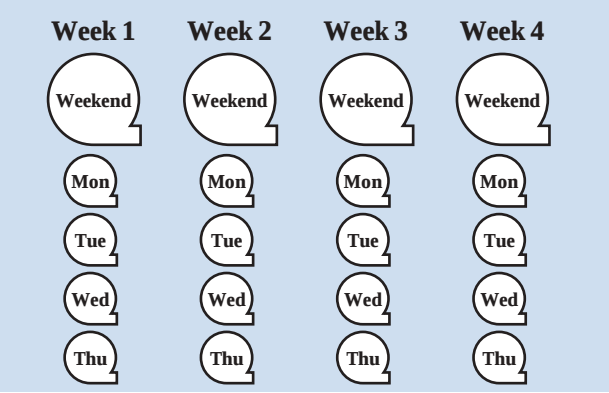

The JCL required to create the backups is very straightforward. Here are two example jobs showing the weekly Full (on the left) and daily Incremental (on the right):

Weekly Full-Volume Backup

```
//BACKUP      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT     DD    SYSOUT=*
//SYSPRIN1     DD    SYSOUT=*
//SYSPRIN2     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=FDR
     MOUNT    VOLG=SYS
     MOUNT    VOLG=PROD
```

Daily Incremental Backup

```
//BACKUP      EXEC   PGM=FDRABR,REGION=0M
//SYSPRINT     DD    SYSOUT=*
//SYSPRIN1     DD    SYSOUT=*
//SYSPRIN2     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
```

In these examples, all volumes currently online which have volsers beginning SYS or PROD will be dumped. The 'DUMP TYPE=FDR' statement (for the full volume backup) and 'DUMP TYPE=ABR' statement (for the incremental) dictate which type of backup will be taken. In both jobs, two tape drives will be used to hold the backups, as requested by the TAPE1 and TAPE2 DD cards. A second copy of each backup will also be stored, as directed by the TAPE11 and TAPE22 DD cards. These secondary copies of the backups will be recorded alongside the primary copies, and the actual tapes can be stored off-site to provide for Disaster Recovery. The primary copies will normally stay on-site to satisfy ad-hoc restore requests. All of the backups that are created will be automatically "piggy-backed", ensuring that every tape is filled to capacity before a new tape is mounted.

PART TWO

DATA BACKUP

Restoring *datasets* from the Volume Incremental backups is also very easy. The sample job below illustrates the restore of several datasets, either from the most recent copy (as in the case of USER.FILE1), or from older copies (USER.FILE2 & USER.FILE3).

Various mechanisms are employed to identify which version of the dataset is required. Notice also on the SELECT statement for the USER.FILE3 dataset that a NEWI operand has been coded to request that the dataset be renamed during the restore with a new index of TEST.

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,NEWI=TEST
```

Restoring a *whole* DASD volume is equally straightforward. With a simple job (see example JCL on the next page), the contents of the disk can be automatically recreated from the daily incrementals and the weekend full, as shown in the diagram here.

The fully automated process ensures that each dataset on the disk is restored once only, and to its original location.

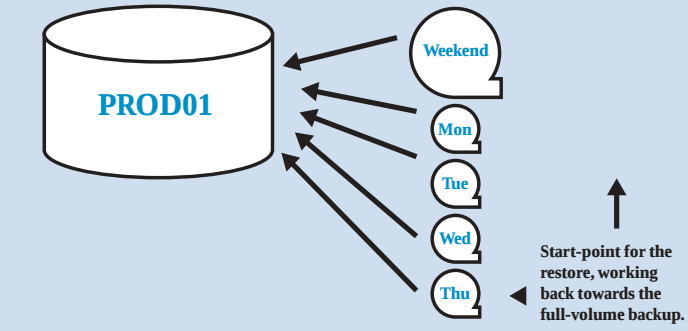

PART TWO

DATA BACKUP

Below is a sample job illustrating the simplicity of the volume restore process. The DASD volume PROD01 is to be recovered to the point of its most recent backup (GEN=CURRENT), overlaying the disk currently online with a volser of SPARE1. The CPYVOLID=YES parameter ensures that the output volume will be labeled as PROD01 after the restore has completed. The DYNTAPE parameter dynamically allocates the tape drive. All tape volumes required for the restore are automatically called for and mounted in the correct sequence.

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,DYNTAPE,CPYVOLID=YES,GEN=CURRENT
    SELECT  VOL=PROD01,NVOL=SPARE1
```

PROD01's most recent full-volume backup is located, together with all the incremental backups that followed it. The most recently created incremental is restored first, rolling back the VTOC, VTOCIX, VVDS and any datasets on that backup. The restore process then works back through the preceding incrementals and the full backup, restoring only the most recent copy of each dataset. The final result is a volume which looks **exactly** like the original volume did at the time of the last incremental. All datasets are in their original locations, with the exact same allocation characteristics.

Multiple DASD volumes can also be restored by this automated process—either one at a time or concurrently. We look at this in more detail in PART THREE, 'Data Recovery'.

"Standard" Volume and Dataset Backups

As well as the automated Volume Incremental backup system described above, it is also possible to create "standard" (*i.e.* manual) backups of datasets and/or DASD Volumes. These backups may be created to cater for unusual ad-hoc circumstances where the automated recording, retention and expiration facilities are not required (*e.g.* when backing up OS/390 and z/OS System volumes). No internal recording mechanism is used for these backups, although they can be cataloged in an MVS ICF user catalog and (optionally) controlled by standard MVS GDG processing.

PART TWO

DATA BACKUP

Application Backup

Although many companies utilize a Volume Incremental system for their day-to-day backups, perhaps complemented by ad-hoc “standard” backups, more and more are also now looking towards Application Backup to meet their specific backup and recovery needs.

The FDRAPPL Application Backup component of the FDR Storage Management System is designed to backup datasets that belong to an application, *regardless* of where these datasets reside.

MVS catalogs are searched to locate the requested datasets, and a record of each one backed up is recorded in a control file, which can be optionally copied to the end of the backup tape.

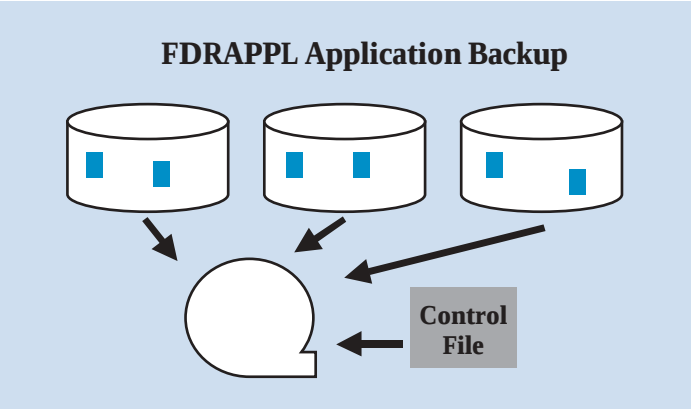

Here is an example job illustrating an FDRAPPL backup of some datasets belonging to a payroll application. The datasets, which are identified by the CATDSN statement, will be backed up and retained for 14 days. Details of the backups will be recorded in the control file pointed to by the ARCHIVE DD statement. Two copies of the backup will be taken simultaneously, to the TAPE1 and TAPE11 DD’s.

Application Backup

```
//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,RETPD=14
    SELECT CATDSN=PAYROLL.**
```

A restore of one or more of the application datasets is equally straightforward, as shown in the sample job below. In this example, all of the datasets belonging to the application are to be recovered (SELECT ALLDSN). Once again, the DYNTAPE parameter will dynamically allocate the appropriate tape drives, with FDRAPPL controlling the requests for the required backup tapes. Information regarding the backups is obtained from the control file which was created during the backup process and which is indicated in the job by the ARCHIVE DD card.

PART TWO

DATA BACKUP

During the restore process, the selected datasets can either be directed back to the DASD volumes from where they were dumped (*i.e.* during a home-site restore of the application), or they can be directed to new DASD volumes in the event of a Disaster Recovery. The RECAT and VRECAT parameters ensure that all the restored datasets are correctly cataloged.

Application Restore

```
//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
```

As you can see from the example jobs, 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. The information stored in the recording mechanism provides for the automated restore processes of individual datasets and/or whole applications, making it an ideal tool for home-site restores or disaster recovery.

The full features of FDRAPPL are described in more detail in the *"FDRAPPL Application Backup"* Concepts and Facilities Guide.

Backup Integration and Performance

All of the backup systems described above use a common 'data mover' module, which is employed by all "data-moving" members of the FDR Storage Management Family.

This **integration** ensures that the control statements used by each of the backup systems have a consistent 'look and feel' about them.

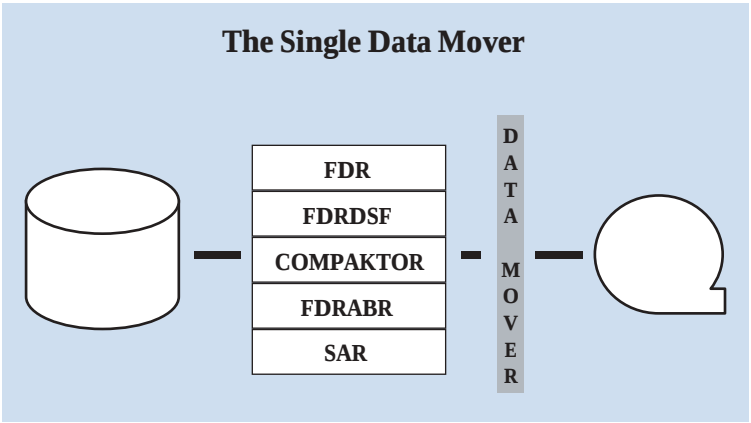

PART TWO

DATA BACKUP

As all the members of the family create 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 (see PART THREE). The employment of the single data mover also provides benefits in **backup performance** due to its economical use of critical resources such as EXCP's and 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:

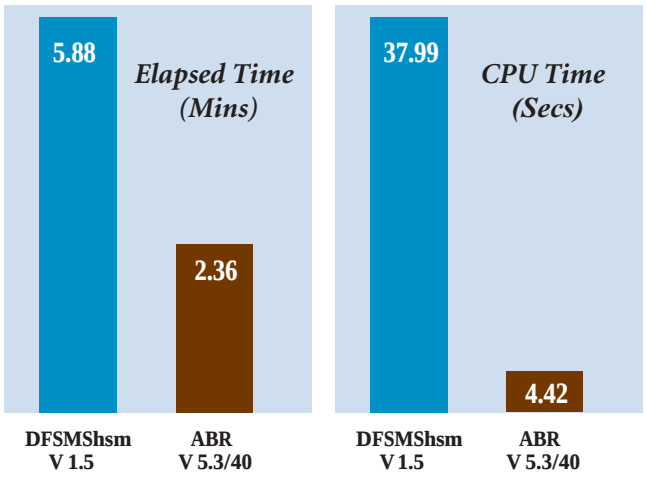

*In this test, an **Incremental Backup** was run against two 3390-2 DASD volumes containing 2,800 datasets, of which 304 were selected for backup (56 VSAM and 248 non-VSAM).*

The backups from the two disks were performed concurrently to two IBM 3590-B11 tape drives on a 3590-A50 controller.

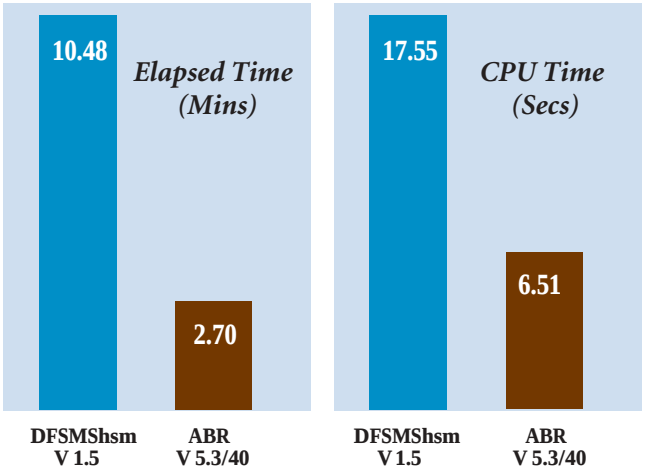

*In this test, an **Application Backup** was run against 1,108 non-VSAM datasets, residing on two 3390-2 DASD volumes on a RAID device.*

The backup was taken to a single IBM 3590-B11 tape drive.

Note: Full details of both of these tests can be obtained on request.

PART TWO

DATA BACKUP

Utilizing FDRINSTANT

Many companies now need to operate 24 x 7 operations and the disruption caused by daily backups, however quick, is no longer tolerable. The performance of the Volume Incremental Backup system and the "Standard" Backup systems described earlier can now be further enhanced by utilizing FDR Instantbackup (FDRINSTANT), which is a separately priced member of the FDR Storage Management Family.

FDRINSTANT works in conjunction with the following hardware-specific products:

- *TimeFinder* - EMC Symmetrix
- *SnapShot Copy* - IBM RVA
- *SnapShot Copy* - StorageTek SVA/V960
- *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 can then be used to take "instant" point-in-time 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:

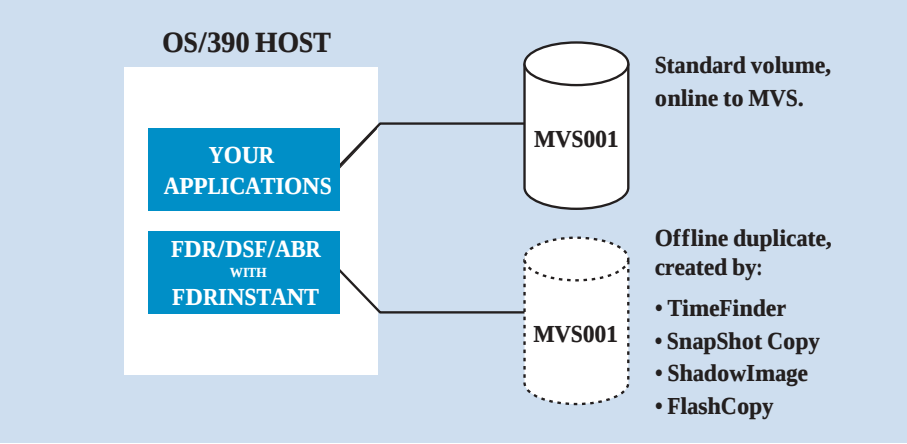

FDRINSTANT does *not* require the duplicate volumes to be brought online. It can access them across the S/390 channel while they remain *offline* to MVS. This eliminates potential problems with duplicate volsers and VTOC/VVDS and Catalog discrepancies.

The Volume Incremental 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, at the time of writing, DFSMShsm has not been similarly enhanced and does not offer an equivalent to FDRINSTANT.

FDRINSTANT also enhances the FASTCPK and FDRCOPY modules of the FDR Storage Management Family (described in PART FOUR—'Storage Management Utilities').

The full facilities of FDRINSTANT are described in the "FDRINSTANT" Concepts & Facilities Guide.

PART TWO

DATA BACKUP

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/V960 and IBM RVA subsystems. HSDM allows disk data to be backed up without first decompressing it.

Depending on the levels of compression achieved by the SVA/V960/RVA, the amount of data that must be read from disk and transferred to the backup media can be reduced by up to 60% when HSDM is employed. 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/V960/RVA, FDRINSTANT also avoids the compression overhead by directly writing the compressed track images. This provides similar reductions in restore elapsed times.

Each of the backup systems described earlier 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 tasks with ExHPDM 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 HSDM and/or 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, a full volume backup was taken 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.

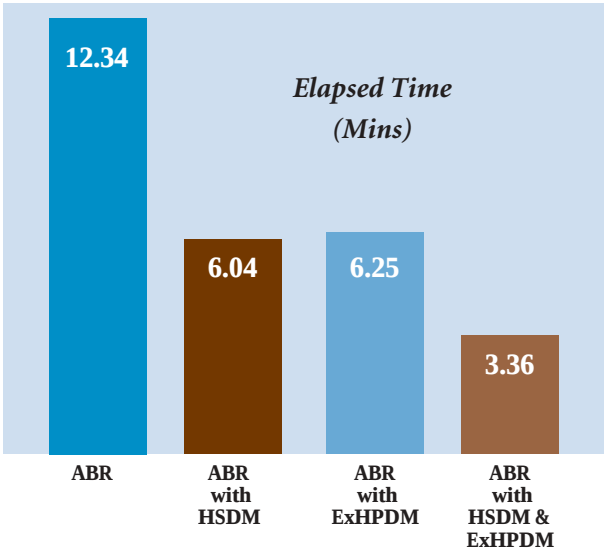

Important Note: FDRINSTANT is required for the support of HSDM and ExHPDM within FDR and ABR, but the above savings in Elapsed Time are possible whether backing up from an offline volume, or from a standard online volume.

PART TWO

DATA BACKUP

Summary

The FDR Storage Management family provides today's OS/390 and z/OS Storage Administrators with a range of backup options.

The **Volume Incremental Backup** system automates the process of taking daily volume-based backups. These backups are fast, efficient, low-resource using processes.

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

The **Standard Backup** tools can be used to satisfy any other ad-hoc backup requirements.

All of the backup systems within the FDR Family use the same high-performance data mover, ensuring a fast and efficient backup and restore, as illustrated by the benchmark figures on page 13.

Furthermore, Innovation is continually enhancing the backup and restore systems 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. This philosophy is perhaps best illustrated by the recent enhancements added by the support in FDRINSTANT for TimeFinder, SnapShot Copy, ShadowImage and FlashCopy. Likewise, the support for HSDM and ExHPDM.

FDR/ABR's Volume Incremental Backup and FDRAPPL Application Backup continue to be the backup systems of choice for Data Centers that are serious about securing their critical corporate data.

PART THREE

DATA RECOVERY

Introduction

PART TWO described the various facilities in the FDR Family for taking volume and/or dataset backups. This flexibility is important for ensuring that corporate critical data is secured in a timely and efficient manner. Of equal importance is the ability to both quickly and easily *recover* data—either with ad-hoc restores at the home-site, or during complete recoveries of applications and volumes in the event of a disaster.

We've already seen from the example restore jobs in PART TWO how easy it is to restore whole DASD volumes (page 10) and individual applications (page 12). Either process may be required to recover from a small-scale problem at the home-site, while a combination of these tasks would probably be employed in a full Disaster Recovery. For now, let's take a closer look at the *speed and efficiency* of the restore tasks.

Restore Performance

We illustrated earlier that the performance of the FDR Family backup systems (*i.e.* "Volume Incremental Backup" and "FDRAPPL Application Backup") are far superior to competing products like IBM's DFSMShsm. The *restore* processes employed by the Innovation systems are also very fast and efficient, as illustrated by the following benchmark results.

A full-volume recovery of two 3390-2 DASD volumes was carried out. A weekend full-volume backup, together with one daily incremental, was used for the recovery.

As you can see, the ABR recovery task used only a fraction of the CPU resources compared to DFSMShsm, and it completed the recovery in less than a third of the time!

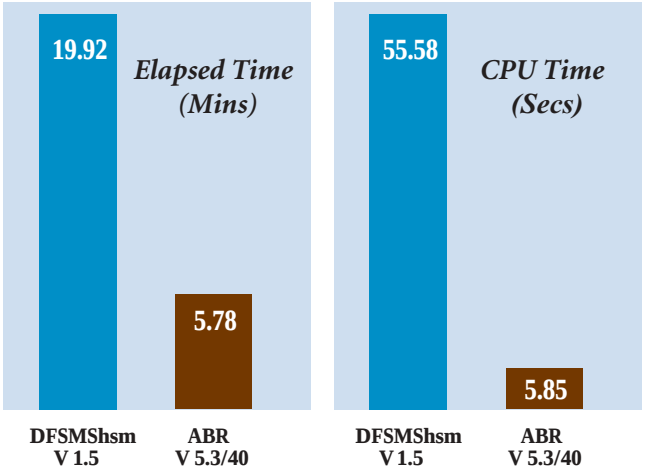

PART THREE

DATA RECOVERY

In another test, the FDRAPPL Application Backup & Recovery system was used to recover the 1,108 non-VSAM datasets that were backed up in the test shown on Page 11.

As you can see, the Innovation restore utility once again out-performed its IBM equivalent in both speed and efficiency.

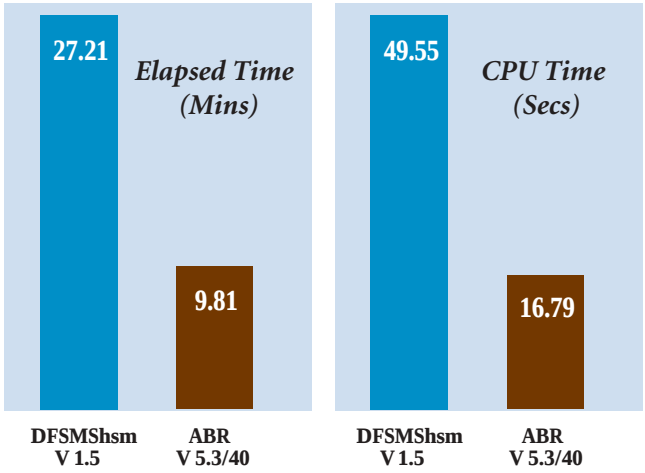

The exceptional performance advantages that the FDR Family enjoys over its competitors should never be underestimated—least of all when looking at Disaster Recovery. The first of the previous two benchmarks illustrates the sheer gulf in performance between an ABR Volume Incremental recovery and the similar (but significantly slower) process employed by DFSMShsm.

Many companies still rely on a volume-based recovery of their production systems during a disaster recovery. For these companies:

- A large number of DASD volumes have to be recovered in one go, often several hundred.
- Each volume must be restored as quickly as possible.
- Each volume must be recovered to its most recent backup (full or incremental).

The speed and efficiency of the ABR Incremental Volume Recovery makes this approach possible. ABR users typically employ *multiple* ABR Volume Recovery jobs, restoring their volumes in parallel, in a fast and efficient manner, and to the point of their most recent backup. This often negates the need to do additional 'application' restores after the volume recoveries have completed.

Quite simply; no other Storage Management system provides the raw performance and efficiency required to recreate a large number of DASD volumes in the time allotted within most Disaster Recovery plans (e.g. 12-24 hours). This is the primary reason why the ABR Volume Incremental Backup system has remained the first choice for Disaster Recovery specialists for over twenty years.

PART THREE

DATA RECOVERY

Utilizing FDRDRP

For optimum performance on a Disaster Recovery (either when testing, or when carrying out a *real* recovery), it is important to restore as many DASD volumes *concurrently* as the number of cartridge drives that have been made available by the DR provider.

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

Although it improves the efficiency and utilization of cartridge space, this consolidation of numerous backup files (from multiple disks) onto a single cartridge can cause serious contention problems when multiple disks need to be restored at the same time. The cartridge containing the backups can only be used by one restore task at a time, so other restores 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 Storage Management Family called FDRDRP (Disaster Recovery Product), which is a *cost-option* to the FDR Storage Management System.

FDRDRP provides an additional layer of automation to the ABR Incremental Volume Recovery described on page 10. By controlling the initiation of *concurrent* volume recoveries, avoiding contention on individual cartridges, and by passing cartridges and drives between the restore tasks, FDRDRP allows many disk volumes to be restored in parallel, ensuring that input cartridges are mounted a minimum number of times—often only once. This can have a dramatic effect on the overall time required to complete the recovery of a large number of volumes. Savings of up to 80% have been reported by some users.

Let's take a look at a simple example to illustrate how FDRDRP operates…

PART THREE

DATA RECOVERY

The diagram below shows the backups of three DASD volumes—PROD01, PROD02 and PROD03. Full-volume backups of these disks were taken at the weekend and written to cartridge 111111. Daily incremental backups were then taken on Monday and Tuesday evenings, written to cartridges 222222 and 333333 respectively.

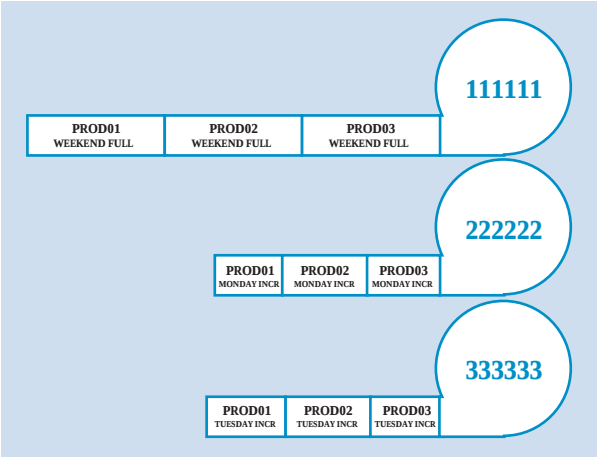

Ordinarily, the standard ABR Full-volume reconstruct process would have to mount and rewind each of the cartridges 3 times (9 mounts/rewinds), taking time to position to the required backup file.

FDRDRP, on the other hand, would mount and rewind each cartridge only *once* (3 mounts/rewinds), thus eliminating all of the positioning delays.

This can result in a significant reduction in the overall time required to recovery a large number of DASD volumes.

At the start of the recovery process, FDRDRP initiates a subtask for each of the three DASD volumes, which are then sorted by the cartridge volser and fileseq required for the first backup. This allows the subtasks to read the backup files on a cartridge in *physical* order with minimal positioning.

PART THREE

DATA RECOVERY

As shown in this table summarizing the restore process, the recovery subtask for PROD01 dynamically allocates and mounts cartridge 333333 and begins the restore from File 1, which is PROD01's Tuesday incremental. The subtasks for PROD02 and PROD03 wait for cartridge 333333.

PROD01	PROD02	PROD03	
Mount tape 333333 Restore from File 1 (Tuesday Incremental)	Wait for tape 333333	Wait for tape 333333	
Mount tape 222222 Restore from File 1 (Monday Incremental)	Pick up tape 333333 Restore from File 2 (Tuesday Incremental)		
Mount tape 111111 Restore from File 1 (Weekend Full)	Pick up tape 222222 Restore from File 2 (Monday Incremental)	Pick up tape 333333 Restore from File 3 (Tuesday Incremental)	Dismount 333333
PROD01 RESTORED	Wait for tape 111111	Pick up tape 222222 Restore from File 3 (Monday Incremental)	Dismount 222222
	Pick up tape 111111 Restore from File 2 (Weekend Full)	Wait for tape 111111	
	PROD02 RESTORED	Pick up tape 111111 Restore from file3 (Weekend Full)	Dismount 111111
		PROD03 RESTORED	

When the PROD01 restore subtask has finished with cartridge 333333, the restore subtask for PROD02 picks it up (without rewinding or dismounting it) and begins the restore from File 2, which is PROD02's Tuesday incremental.

The restore subtask for PROD01 then mounts cartridge 222222 and starts the restore from the Monday incremental.

And so the process continues…until all the weekend Full backups have been restored and all the cartridges dismounted.

This was a simplified example of the FDRDRP-driven restore process. In reality, tens or maybe even hundreds of DASD volumes will need to be restored to complete a full disaster recovery. Most DR providers will have a large number of cartridge drives available and it is *essential* that these drives are used as effectively as possible—within the constraints of the increased capacity of the backup cartridges.

As a slightly expanded example, the sample FDRDRP job below shows the concurrent recovery of 10 DASD volumes, as indicated by the SELECT statements. Since the volumes at the disaster site are pre-initialized to a known volser, each SELECT specifies 'VOL=' for the volume to be restored and 'NVOL=' for the target volume to restore to. After the restore, each target volume will be relabeled to the serial of the restored disk.

PART THREE

DATA RECOVERY

As this is a recovery at a DR site, the COPY=2 parameter is coded to tell FDRDRP to locate the most-recently created COPY 2 backups for each disk volume. No more than 8 cartridge drives will be used, and FDRDRP will mount the backup cartridges the minimum number of times necessary to restore data for all disks. All required disks and tapes will be dynamically allocated.

FDRDRP Restore

```
//RESTFULL EXEC PGM=FDRDRP,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSIN     DD  *

   RESTORE TYPE=DRP,MAXTAPES=8,CPYVOLID=YES,COPY=2
   SELECT  VOL=PROD01,NVOL=DR0170
   SELECT  VOL=PROD02,NVOL=DR0171
   SELECT  VOL=PROD03,NVOL=DR0172
   SELECT  VOL=PROD04,NVOL=DR0173
   SELECT  VOL=PROD05,NVOL=DR0174
   SELECT  VOL=PROD06,NVOL=DR0175
   SELECT  VOL=PROD07,NVOL=DR0176
   SELECT  VOL=PROD08,NVOL=DR0177
   SELECT  VOL=PROD09,NVOL=DR0178
   SELECT  VOL=PROD10,NVOL=DR0179
```

An FDRDRP job can use multiple cartridge drives, and multiple concurrent FDRDRP restore jobs can be initiated, each one restoring a different set of disk volumes. If two jobs require the same cartridge at the same time, FDRDRP will pass the cartridge from one job to the other without dismounting or rewinding it.

With this highly efficient utilization of cartridge drives, FDRDRP offers the power and control required to get the very best out of the facilities provided. Even if you are only *testing* the data recovery process (*i.e.* you are not doing a real disaster recovery), it is vitally important to complete the recovery in the shortest possible time. This reduces costs and also allows the end users the maximum amount of time to actually test the recovered systems. Of course, on a *real* disaster, the time saved by FDRDRP could mean a lot more than that.

Summary

The FDR Storage Management Family provides today's OS/390 and z/OS Storage Administrators with all the tools required to recover their data as quickly as possible—either for ad-hoc restores at the home-site or for full-blown Disaster Recovery.

As shown by the benchmarks on pages 17and 18, even the closest competitor cannot provide anything like the restore performance offered by FDR/ABR. With the addition of the FDRDRP module, the FDR Family's 30-year reign as the Number One Disaster Recovery tool looks set to continue for some time to come.

PART FOUR

STORAGE MANAGEMENT UTILITIES

Introduction

The previous sections of this Guide have concentrated on the backup and restore facilities provided by the FDR Family. However, Storage Management is about more than just taking backups, so the Family also includes a range of tools and utilities to help today's Storage Managers get the very best from their DASD. In summary, those tools are:

- FDRABR - Data Migration & Recall
- FDRCOPY - Dataset Copy/Move
- FASTCPK - DASD Volume Reorganization
- FDRREORG - VSAM and IAM Reorganization and PDS Compression
- FDRPAS - Non-Disruptive Movement of DASD Volumes

As well as forming part of the FDR Storage Management Family, some of the above modules are also available as separately priced options (contact Innovation for details). Let's take a look at each one now.

FDRABR—Data Migration & Recall

Although the price per megabyte of OS/390 disk storage has dropped considerably over the past few years, the demand has increased in most installations. For this reason, disk storage remains a significant part of the Data Center budget.

The available DASD space at an installation, however cheap, needs to be used as efficiently as possible. Unfortunately, DASD volumes often contain datasets that no longer need to be on disk. In many cases, datasets are created by a batch job or TSO user and never used again. Other datasets are used only on an infrequent basis—perhaps just once a year—and do not need to be on disk in the interim period. In addition, there are also the datasets that do not meet installation standards (*e.g.* uncataloged/expired) which need to be regularly dealt with.

The FDRABR Data Migration & Recall System provides an excellent way to control the consumption of DASD resources by ensuring that expired or unwanted datasets are either deleted or moved to a less expensive medium (*e.g.* tape or compressed DASD).

Here is a sample job illustrating the ease-of-use of the FDRABR Migration System.

Datasets which have not been referenced in the last 60 days (indicated by the ADAYS=60 keyword) will be migrated from all online volumes.

FDRABR Migration—Example 1

```
//ARCHIVE   EXEC PGM=FDRABR,REGION=0M
//SYSPRINT    DD  SYSOUT=*
//SYSPRIN1    DD  SYSOUT=*
//ARCHIVE     DD  DSN=FDRABR.ARCHIVE,DISP=SHR
//TAPE1       DD  DSN=ARC1,DISP=(,KEEP),UNIT=TAPE
//SYSIN       DD  *

    DUMP    TYPE=ARC,ONLINE,ADAYS=60,RETPD=730,DSNENQ=USE,RECALL=YES,MIGRAT=YES
```

PART FOUR

STORAGE MANAGEMENT UTILITIES

After being deleted from disk, the datasets will be recorded in an Archive Control File, which is a very compact/efficient file that can hold the details of several hundred datasets in a single track. The datasets will be recataloged for Auto-Recall (RECALL=YES), allowing batch jobs and online users to perform quick and transparent recalls if required, without causing unwanted delay or disruption. The MIGRAT=YES option ensures that migrated datasets are cataloged to MIGRAT, which is recognized by IBM as indicating a migrated dataset. If the datasets are not recalled after 2 years and one month (RETPD=730), they will be removed from the control file.

For Data Centers without automated tape libraries (ATLs), or for those that require the fastest possible recall, the migrated data can be optionally directed to disk. The example below illustrates the use of DASD *and* Tape to hold the migrated data. The first copy, which will be retained for 30 days, will be sent to a pool of disks, as directed by the TAPE1 DD statement. Any initial recalls (*i.e.* within the first 30 days) can be satisfied from disk without involving a tape mount. A second copy of the data will also be sent to tape (via the TAPE11 DD) and will be retained for 730 days.

FDRABR Migration—Example 2

```
//ARCHIVE  EXEC PGM=FDRABR,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSPRIN1  DD  SYSOUT=*
//ARCHIVE   DD  DSN=FDRABR.ARCHIVE,DISP=SHR
//TAPE1     DD  DSN=FDRABR.POOLDISK.POOL1,UNIT=DISK,
//            VOL=SER=(ARC001,ARC002,ARC003),DISP=OLD,LABEL=RETPD=30
//TAPE11    DD  DSN=SMS2,UNIT=TAPE,DISP=(,KEEP),LABEL=RETPD=730
//SYSIN     DD  *

    SIM     TYPE=ARC,SMSMANAGE=YES,RECALL=YES,DSNENQ=USE,MIGRAT=YES
    MOUNT   STORGRP=TSO1
    MOUNT   STORGRP=TSO2
```

The Data Migration in the above example is being limited to datasets residing on DASD volumes currently belonging to the TSO1 and TSO2 DFSMS Storage Groups, as directed by the MOUNT STORGRP statements. When operating against DFSMS-controlled DASD Volumes, either ABR selection criteria can be used, or (if SMSMANAGE=YES is coded, as above) the attributes of the SMS Management Class associated with each dataset can be employed in the selection process. If the Management Class is chosen, and both SMS and non-SMS volumes are processed in the same run, ABR criteria will be used for the non-SMS volumes and ignored for SMS volumes.

Notice also that the above uses a “SIM TYPE=ARC” control statement instead of the “DUMP TYPE=ARC” utilized in Example 1. This is an illustration of the simulation feature available with FDRABR Migration—an essential tool for ensuring that the selection criteria will do what you expect it to do *before* you run the job for real!

Additional ‘pre-migration’ reporting is also provided by the FDRQUERY reporting tool, as described in PART FIVE.

PART FOUR

STORAGE MANAGEMENT UTILITIES

FDRCOPY—Dataset Copy/Move

FDRCOPY is a utility included in the FDR Storage Management Family for copying or moving datasets from one DASD volume to another. FDRCOPY can be used in place of IEBCOPY, either for day-to-day use, or for special DASD migrations or DFSMS conversions.

- Most types of dataset are supported, including Sequential, PDS, PDSE, VSAM, DB2, Striped, etc.
- FDRCOPY can copy/move single datasets, groups of datasets, or all datasets on a volume. An unlimited number of datasets can be copied in one execution.
- Datasets can be copied or moved to the same name, or to a new name. They can be copied or moved from any number of input disk volumes, and can be directed to any number of output disk volumes.

FDRCOPY can also invoke FDRREORG (see FDRREORG section later) to compress PDS datasets.

Below is an example of an FDRCOPY ***copy*** job to copy and rename a group of VSAM (DSORG=EF) datasets which have been located via the MVS ICF user catalogs. All datasets beginning with 'ABC' will be copied across to new volumes (PROD01 and PROD02) and the copies will be given a new first level index of 'XYZ'. FDRCOPY will default to cataloging the newly created datasets in the appropriate user catalog.

FDRCOPY—COPY Example

```
//COPY     EXEC PGM=FDRCOPY,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSIN     DD  *

    COPY   TYPE=DSF
    SELECT CATDSN=ABC.**,NEWI=XYZ,DSORG=EF,NVOL=(PROD01,PROD02)
```

The next example shows an FDRCOPY ***move*** job to consolidate datasets scattered across a group of volumes (DEV) onto a new single volume (TST001). The catalog entry for each dataset that is moved will be updated to point to TST001.

FDRCOPY—MOVE Example

```
//MOVE     EXEC PGM=FDRCOPY,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSIN     DD  *

    MOVE   TYPE=DSF
    SELECT ALLDSN,VOL=DEV*,NVOL=TST001
```

PART FOUR

STORAGE MANAGEMENT UTILITIES

FDRCOPY with FDRINSTANT

When used in conjunction with FDRINSTANT, FDRCOPY can take advantage of the special hardware features available in the IBM RVA and StorageTek SVA/V960 (with SnapShot Copy), and the EMC Symmetrix (with TimeFinder).

FDRINSTANT enhances FDRCOPY to copy the required disk tracks *within* the disk subsystem, without transferring data up and down the disk channels.

As shown in the diagram, if the input and output disks are in the same disk subsystem, FDRCOPY can instruct the disk subsystem to copy the tracks *internally*.

This feature removes the overhead of the channel transfers and *significantly* improves the performance of the FDRCOPY task.

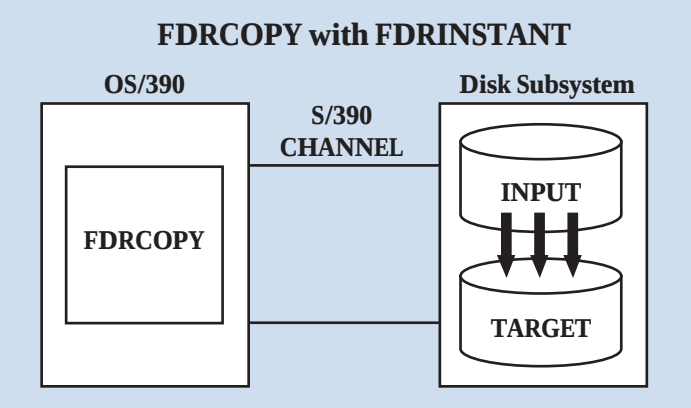

With an 'instant' FDRCOPY, datasets can be copied in a fraction of the time it takes with 'standard' FDRCOPY. No special parameters have to be coded; FDRCOPY will automatically detect when the input and output disks are in the same disk subsystem. If multiple datasets are being copied, FDRCOPY will use the 'instant' method wherever possible, resorting to the 'standard' method only for the copying of datasets across different disk subsystems.

In this next example, some 'PROD' datasets are being copied across to two other online volumes (DEV001 and DEV002), which are within the same SVA/V960 or RVA subsystem. A new second-level index of 'TEST' is being inserted in the names of the copied datasets. As the input and output volumes are within the same Storage subsystem, FDRCOPY automatically invokes the SnapShot hardware to do an instantaneous copy.

'Instant' Copy datasets within the same SVA/RVA.

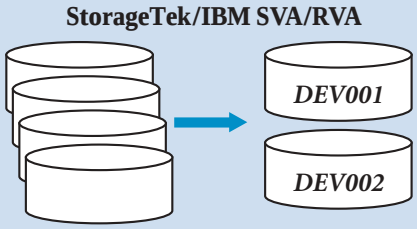


```
//COPYFILE  EXEC PGM=FDRCOPY,REGION=0M
//SYSPRINT   DD  SYSOUT=*
//SYSIN      DD  *

    COPY    TYPE=DSF
    SELECT  CATDSN=PROD.**,NEWI=.+TEST,NVOL=(DEV001,DEV002)
```

Note: For users with multiple types of DASD Subsystem (and the appropriate FDRINSTANT licensing), the same FDRCOPY job can be used to do ultra-fast copies/moves of datasets, regardless of the platform in use.

STORAGE MANAGEMENT UTILITIES

FDRINSTANT also enhances FDRCOPY to use the *offline* volumes created by SnapShot Copy and TimeFinder. In this final example, DASD volumes SMS001 and SMS002 have been assigned BCV's at addresses 01F4 and 01FA. The user wants to copy some datasets for parallel batch processing, but can only quiesce updates for a few moments:

Step 1: Split the BCV's

```
//STEP1    EXEC PGM=EMCTF
//SYSOUT    DD  SYSOUT=*
//SYSIN     DD  *

     SPLIT 1,01F4,WAIT
     SPLIT 1,01FA,WAIT
```

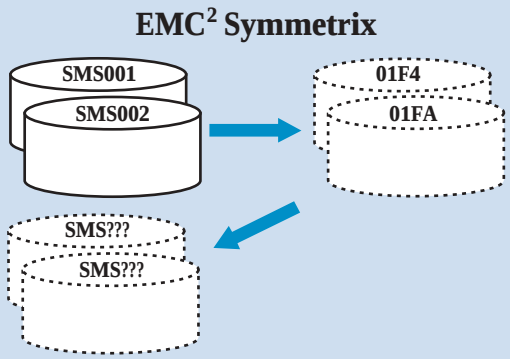

Step 2: Copy across to online volumes

```
//STEP2    EXEC PGM=FDRCOPY,REGION=0M,COND=(0,NE,SPLIT)
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=*
//DISK1     DD  DSN=FDR.USE.UNIT01F4,DISP=OLD
//DISK2     DD  DSN=FDR.USE.UNIT01FA,DISP=OLD
//SYSIN     DD  *

     COPY TYPE=DSF,BCV=(USE,RET)
     SELECT DSN=PROD.PAYROLL.**,NEWI=..+COPY,STORCLAS=BATCH
```

- Step 1 invokes the EMCTF program to split the BCV's from their paired volumes, creating the point-in-time images. The program waits for the split to complete.
- Step 2 then invokes FDRCOPY to copy the selected datasets from the offline BCV's across to other online volumes, renaming the datasets during the copy process.

After the copy is complete, FDRCOPY will then re-synchronize the BCVs with their online paired volumes.

FASTCPK—DASD Volume Reorganization

Modern virtual disk arrays, such as the IBM ESS Shark and the StorageTek SVA/V960, are still regarded by the OS/390 operating system as a set of conventional DASD volumes. Each logical volume in the array has a fixed number of cylinders, and space is managed in the usual way—with a standard VTOC and VTOCIX.

If datasets are over-allocated, then the unused space on the volume is not available for allocation to other datasets. The total number of extents per volume for the various dataset types is limited (*e.g.* for some, the limit is still just 16 extents per volume) and if freespace on a volume becomes fragmented, it can be difficult to allocate new datasets. Additional logical volumes can be defined in the DASD Subsystem to provide extra space, but this adds to administrative and processing overheads, and there are limitations on the total number of volumes (addresses) that can be defined within the array.

PART FOUR

STORAGE MANAGEMENT UTILITIES

FASTCPK is designed to overcome the above problems, with key features including:

- Freespace Fragmentation—FASTCPK can consolidate the freespace on a volume into as few as one or two contiguous areas.
- Multiple Extent Consolidation—FASTCPK can merge the extents of Sequential, PDS, PDSE, VSAM, DB2 and Extended Format datasets—obviating the need to run DFSMShsm extent consolidation.
- Space Release—FASTCPK can release all or part of the unused space within Sequential, PDS, PDSE, VSAM, DB2 and Extended Format datasets.
- VTOC manipulation—FASTCPK can perform a complete analysis of a VTOC, detecting (and in most cases fixing) logical errors, such as invalid Format 5 DSCB's and incorrect freespace definitions. Volumes can be selected by specific or masked volsers, or by SMS Storage Group names.

FASTCPK is ultra-fast! Instead of chipping away at a volume over several hours, it can get the job done in just a few minutes. For example, when run against a 3390-3 in a virtual disk array, the average time required to consolidate free space, merge extents, and release unused space is about 2 minutes. If you just want to release unused space, the TYPE=RLSE mode usually runs in *less than 10 seconds!*

The example job below illustrates a FASTCPK against a group of TSO and TEST DASD volumes, identified by the 'VOL=' parameter. The Compaktion will only take place on each volume if the number of freespace extents exceeds 20, as dictated by the CPKFREEX parameter. (The IBM Fragmentation Index can also be used to identify and select badly fragmented volumes). The PSRLSE, PORLSE and VSRLSE parameters specify that all unused space will be released from Sequential, PDS, PDSE and VSAM datasets—but only if they have a secondary allocation coded. (NOSECOND=NORLSE).

FASTCPK Example

```
//FASTCPK  EXEC PGM=FDRCPK,REGION=0M,COND=EVEN
//SYSPRINT  DD  SYSOUT=*
//SYSMAP    DD  SYSOUT=*
//SYSSUMM   DD  SYSOUT=*
//SYSIN     DD  *

    COMPAKT TYPE=FASTCPK,VOL=(TSO*,TEST*),CPKFREEX=20
    NOSECOND=NORLSE,PSRLSE=ALL,PORLSE=ALL,VSRLSE=ALL
```

PART FOUR

STORAGE MANAGEMENT UTILITIES

FASTCPK can also produce a range of reports, in both detailed and summary format. The MAP function allows you to map out the contents of selected volumes, while the SIMULATE feature lets you see the effects of a Compaktion without actually running it for real. *'Before'* and *'After'* maps can be produced for each volume processed, either on a SIMULATE or a real Compaktion, together with a concise summary of the effects of the Compaktion. Here is an example of the summary report:

```
CPK301I INNOVATION DATA PROCESSING - COMPAKTOR VER. 5.3/01P COMBINED SUMMARY DATE 94.145 TIME 11.32.16 PAGE 1

               - NUMBER OF -    >1  -- ALLOCATED ---  ------- FREE -------   FRAG - EMPTY TRACKS IN -   - VTOC -   TIME COMP
VOLSER DEVTYPE   TRACKS  DSNS  EXTS TRACKS   EXTS %AL TRACKS AREAS LARGEST  INDEX   VSAM    PS     PO   SIZE %US  (MIN) CODE

IDPLB0 3380      13275   104     9   8084    132  61   5191    31    2651  0.232     16     2   1367    15  15
       --AFTER-CPK-->    104     0   8084    105  61   5191     1    5191  0.000     16     2   1367    15  15    1.1    0

TSOWK1 3380      13275   249     7   7591    268  57   5684    31    2205  0.178      0   119    268    10  48
       --AFTER-CPK-->    249     0   7482    141  56   5793     2    2909  0.090      0   119    268    10  48    1.6    0
```

The report shows the state of each volume before and after the real or simulated Compaktion, including the number of tracks allocated, the number of extents on the volume and the percentage allocation of the volume. It also shows the number of freespace areas and the size of the largest free area. The IBM fragmentation index is also shown. In this particular run, no Space Release parameters were coded, so the empty tracks reported in VSAM/PO/PS datasets have not been released. On the far right of the report is the time (real or estimated) for the Compaktion of that volume.

FASTCPK with FDRINSTANT

As with FDRCOPY described earlier, FASTCPK can also take advantage of the facilities provided by FDRINSTANT.

When used in conjunction with FDRINSTANT, FASTCPK's performance can be further enhanced by interacting with IBM/StorageTek SnapShot Copy and EMC's TimeFinder.

As shown in the diagram, FDRINSTANT enhances FASTCPK to move the disk tracks *within* the disk subsystem, without transferring tracks up and down the disk channel. This drastically reduces the overall time required to complete the Compaktion.

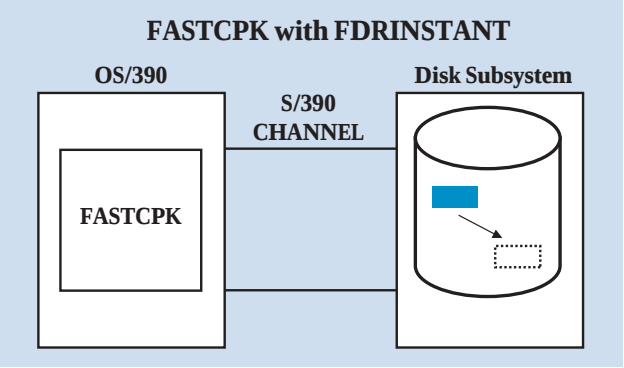

The example below illustrates FASTCPK using FDRINSTANT (via the SNAPSHOT=YES keyword) to interact with SnapShot Copy. Typically, an “Instant” Compaktion against a 3390-3 in an RVA or SVA/V960 with SnapShot Copy will run in *less than 10 seconds!*

FASTCPK—FDRINSTANT Example

```
//FASTCPK  EXEC PGM=FDRCPK,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//SYSMAP    DD  SYSOUT=*
//SYSSUMM   DD  SYSOUT=*
//SYSIN     DD  *

    COMPAKT TYPE=FASTCPK,STORGRP=PROD,SNAPSHOT=YES
```

FDRREORG—VSAM, IAM and PDS Reorganization

FDRREORG is a separately priced option in the FDR Storage Management Family, designed to logically reorganize VSAM KSDS and AIX datasets. It can also reorganize IAM files (see the ‘IAM’ Concepts & Facilities Guide for more information) and it can compress partitioned datasets (PDSs). The reorganization process is based on user-specified exclusion and selection criteria, allowing datasets to be reorganized on an as-needed basis. FDRREORG can also be run in simulation mode to obtain a report of those datasets meeting the user’s selection criteria.

If FDRREORG is unable to obtain exclusive control of a dataset, the dataset will be bypassed and FDRREORG can be instructed to retry at a later point by issuing the appropriate system ENQs to process the dataset when it becomes available. If it is critical that a particular dataset be reorganized, FDRREORG can be instructed to wait for that dataset to become available.

All IAM or VSAM datasets selected for reorganization are backed up to either Tape or DASD, and then immediately reloaded from the backup.

The backup files created by FDRREORG are logical backups obtained by using standard access method interfaces. They can be either normal sequential datasets, or they can be set up as GDG’s. If the GDG format is chosen, FDRREORG will dynamically define the base GDG to an ICF catalog if one does not already exist.

The names of the backup datasets are generated dynamically from the target dataset names, coupled with the user-specified BACKUPINDEX parameter to insert/delete index levels.

All backup datasets are cataloged after they have been created and will be uncataloged and deleted after the reload has completed (unless the backups are to be kept).

FDRREORG provides ‘last tape’ support on a jobname basis, allowing users to append new backups to the last tape used in the previous execution of the same job.

In order to minimize the manual effort required to recover from a failure during FDRREORG processing, a checkpoint and log file are also created to record information that simplifies the recovery of a failed reorganization.

PART FOUR

STORAGE MANAGEMENT UTILITIES

In this first example job, all VSAM datasets (DSTYPE=VSAM) with a high-level qualifier of 'Account' will be located via the catalog (SELECT CATDSN=ACCOUNT) and will be reviewed by FDRREORG. Only those datasets which have more than 5 CA Splits and/or more than 10 CI Splits will be reorganized. The backup files will be GDGs (BACKUP=GDG) and they will have the same name as the original files, with 'BACKUP' added as a suffix.

FDRREORG—VSAM Example

```
//REORG    EXEC PGM=FDRREORG,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//REORGPRT  DD  SYSOUT=*
//REORGRPT  DD  SYSOUT=*
//SYSIN     DD  *

    REORG  BACKUP=GDG,BACKUPUNIT=3480,BACKUPINDEX=++BACKUP
    SELECT CATDSN=ACCOUNT.**,DSTYPE=VSAM,IFANY,CASPLITR>5,CISPLITR>10
```

After an execution of FDRREORG, a summary report is produced (written REORGRPT DD) to show which files were selected for processing, together with information about their before/after status.

Here is a sample of the summary report showing an FDRREORG against some VSAM files.

```
REORG BACKUP=GDG
SELECT CATDSN=(PROD.**,SMPE.**),DSTYPE=VSAM,CASPLITR>4,CISPLITR>4

                                          CISR  CASR                                RECORDS
                                          %OFU  %PEU  TRKS   TRKS  PCT  NUM     LOADED     TRKS  TRKS    PCT  NUM
DATASET NAME              VOLUME   AM  STATUS     DIRU  DIRA  USED   ALOC  USE  EXT   DIRU  DIRA  USED  ALOC    USE  EXT
------------------------  ------   --  ---------- ----  ---   ---    ---   ---  ---   ----  ----  ----  ----    ---  ---
PROD.ACCTPAY.MASTER       PROD01   VS  REORGANIZED  13    27    55     55   100   10    5125        45    55     82   10
PROD.ACCTPAY.ALTMASR      PROD01   VS  REORGANIZED  26    53    85     85   100   16    7367        50    85     59   16
PROD.HRIS.EMPLOYEE.MASTER PROD02   VS  REORGANIZED  49    47    95     95   100   18   10137        70    95     74   18
PROD.PAYROLL.MASTER       PROD03   VS  REORGANIZED  13    27    55     55   100   10    5226        45    55     82   10
PROD.PAYROLL.SUSPENSE     PROD03   VS  REORGANIZED  23    39    65     65   100   12    6124        45    65     69   12
SMPE.SP223.CSI            DLIB01   VS  REORGANIZED   5    16  1305   1350    97    1   99200      1125  1350     83    1
SMPE.SP422.CSI            DLIB01   VS  REORGANIZED   2     6   780    900    87    1   65222       735   900     82    1
SMPE.GLOBAL.CSI           DLIB01   VS  REORGANIZED  56    70   300    330    91    1   21805       150   330     45    1
```

When processing a large number of datasets, FDRREORG provides a multi-tasking capability that allows multiple volumes to be processed in one execution. Each volume selected for processing will be processed by a separate subtask. Up to 15 concurrent subtasks can be active, provided there is sufficient virtual storage available within the region.

STORAGE MANAGEMENT UTILITIES

In this next example, IAM compatible format files beginning with 'Payroll' will be reorganized if the amount of storage required to hold the Overflow Index is greater than 1Mb. The backups, which will be dynamically allocated, will be written to disk. Since there are a large number of IAM files belonging to the PAYROLL application, and these files are spread across a number of DASD volumes, up to 8 subtasks (MAXTASKS=8) will be used to complete the reorganization process.

FDRREORG—IAM Example

```
//REORGIAM EXEC PGM=FDRREORG,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//REORGPRT  DD  SYSOUT=*
//REORGRPT  DD  *
//SYSIN     DD  *

    REORG  MAXTASKS=8,BACKUPUNIT=DISK
    SELECT CATDSN=PAYROLL.**,DSTYPE=IAM,OVERFLOWINDEX>=1048576
```

FDRREORG can also be used to reorganize PDSs, compressing all active members to the beginning of the PDS to allow room for new members to be added. This is similar to an IEBCOPY compress-in-place, but FDRREORG is 50% to 90% faster than an equivalent IEBCOPY, with similar reductions in CPU and I/O resources.

In this final example, FDRREORG is being run against all TSO volumes (VOL=TSO*), searching for PDS datasets which are more than 80% full.

FDRREORG—PDS Example

```
//REORGPDS EXEC PGM=FDRREORG,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//REORGPRT  DD  SYSOUT=*
//REORGRPT  DD  *
//SYSIN     DD  *

    REORG
    SELECT ALLDSN,VOL=TSO*,DSTYPE=PDS,PDSFULL>80
```

Note: The FDRCOPY component of the FDR Storage Management Family (described earlier) can optionally invoke FDRREORG to compress PDSs. This facility may be automatically invoked, allowing user requests for IEBCOPY in-place compressions to be routed through to FDRREORG to take advantage of its superior performance.

PART FOUR

STORAGE MANAGEMENT UTILITIES

FDRPAS—Non-Disruptive Movement of Disk Volumes

FDRPAS is the latest addition to the FDR Storage Management Family. It is available as a cost option and it provides the ability to do ***non-disruptive movement*** of OS/390 disk volumes from one disk device to another. FDRPAS can swap volumes within a single system image, or volumes attached to multiple systems or LPARs in a shared-DASD complex or sysplex, whether locally or remotely attached. Multiple volumes can be swapped concurrently.

FDRPAS supports a wide variety of disk devices from hardware vendors including IBM, EMC, StorageTek, Amdahl and Hitachi. It can swap disk volumes between disks of the same type from the same hardware vendor, or between disks supplied by different vendors. No special software or hardware modifications are required.

FDRPAS volume movement can be done during normal system operations, without interrupting other activity. The operating system, application jobs, online systems, and end-users are unaware that FDRPAS is swapping disks to new devices. Because of this non-disruptive nature, FDRPAS has several distinct uses for today's OS/390 Storage Administrators:

- Installing New DASD Hardware

  OS/390 hardware and software allows you to attach new disk subsystems and dynamically activate an updated I/O configuration to make them immediately available.

  FDRPAS complements that capability by allowing you to move the data from existing volumes to the new hardware while those volumes are still in use. When FDRPAS completes the swap of a disk, the volume resides completely on the new device, and the original device is no longer required.

- General Data Movement

  Aside from when installing new hardware, it is often a requirement to move data around, either within a single DASD subsystem, or across subsystems. For companies who run 24 x 7 operations, and perhaps even installations with large maintenance windows, there simply isn't sufficient 'quiet' time available to move what is often now terabytes of data.

  FDRPAS removes this problem by allowing you to non-disruptively move terabytes of data during normal business hours.

- I/O Load Balancing

  Occasionally, the I/O load in one or more disk subsystems is excessive, while other subsystems remain under utilized.

  FDRPAS can be used to rearrange the disk volumes to distribute the I/O load and improve performance. This work can be done as soon as it is required—day or night.

- Non-Disruptive Backup

  As described in PART TWO of this guide, today's 24 x 7 shops cannot afford the downtime caused by taking daily backups—even when high-speed backups (like those provided by FDR) are employed.

  FDRPAS can be used to create offline point-in-time copies of DASD volumes, which can be used by FDR in conjunction with FDRINSTANT.

STORAGE MANAGEMENT UTILITIES

The process employed by FDRPAS is very simple:

- A monitor job/task is started on each system that has access to the volume(s) to be swapped.
- A simple FDRPAS job is run, or a console command is issued (see below), to request that an online volume be swapped to an offline disk device. The initiation can be done on any system in the complex, and the other systems will automatically join in the swap operation.
- FDRPAS will verify that the target device is offline to all sharing systems, to ensure that an active volume cannot be accidentally overlaid.

Here is an example of the job required to initiate the SWAP:

FDRREORG—SWAP Example

```
//SWAP      EXEC  PGM=FDRPAS,REGION=0M
//SYSPRINT   DD   SYSOUT=*
//SYSIN      DD   *


     SWAP    TYPE=FULL
     MOUNT   VOL=TSO001,SWAPUNIT=2C31
```

Alternatively, on a single system, the following console command can be issued:

```
S PASPROC.TSO001,PARM='SWAP TYPE=FULL/MOUNT VOL=TSO001,SWAPUNIT=2C31'
```

A set of ISPF panels is then provided for monitoring and controlling all of the FDRPAS swap operations that are in progress:

```
──────────────────────────── FDRPAS Plug & Swap ──────────────────────────── Row 1 to 1 of 1
 COMMAND ===>                                                  SCROLL ===> PAGE
Valid commands are: Active, COnfirm, SWap, DUmp, ABort, OPtions, HIstory
Command  Volume   Unit   Swap to                                          REfresh 0
          Serial  Addr   Offline
           Mask    Mask    Unit       Status
_______  ______  ____  ______   ____________________________________________________
          TSO001  07C1    2C31        ACTIVE (MAIN)
          Pass: 1     45 % Tracks to copy: 17120      Copied: 7704      Updated: 15
          Source - Reserve: 0    Level: 1    Pace: 0
          Target - Reserve: 1    Level: 1    Pace: 10
```

For each requested volume, FDRPAS will copy all allocated tracks to the new disk device, while simultaneously detecting updates to the original volume. Updated tracks are re-copied if necessary.

The new device remains offline during the copy process, so that the copied data is protected until the swap is complete. Only the source and target devices are accessed by FDRPAS. It does not use any additional communication between systems and it does not require TCP/IP, VTAM, a dataset on a third disk volume or a coupling facility. A swap of a volume can be terminated at any time before the final swap without affecting the original device or any applications using it.

The swap is accomplished with minimal impact on the performance of other applications using the volumes being swapped, which continue to execute, unaware that the data movement is occurring. FDRPAS dynamically manages the copy process in response to system activity (*e.g.* copying inactive datasets before active datasets and pacing the copy I/O), to further minimize its effect on the system. A swap will normally take between 6-12 minutes for a 3390-3 DASD volume.

At the successful completion of the copy operation, FDRPAS switches devices so that all I/O is directed to the new device. During this process, the new device is placed online, while the old device is varied offline and the label is modified so that OS/390 will no longer be able to vary it online. When the system is re-IPLed, it will be the new device that automatically comes online.

When the SWAP operation has terminated, FDRPAS can, if desired, send an e-mail message to one or more recipients informing them of the successful (or unsuccessful) completion of the task. This feature is enabled via a simple DD statement in the executing JCL, which points to control statements defining the message content and the intended recipients.

As mentioned earlier, after a successful swap, the now-offline original device can be used by FDR and FDRINSTANT as input to a non-disruptive, point-in-time backup of the volume.

In addition to standard 'data' volumes, FDRPAS can swap any OS/390 or z/OS system volume, including the SYSRES and other volumes containing open catalogs. SMS-managed volumes and work volumes are also supported. FDRPAS can even swap a smaller device to a larger device of the same type (*e.g.* 3390-3 to 3390-9), during which process it automatically updates the VTOC and Indexed VTOC on all sharing systems. The only volumes not supported by FDRPAS are those containing active PAGE or SWAP datasets.

Independent of SWAP operations, FDRPAS also includes a utility function (SWAPBUILDIX) which allows you to build an indexed VTOC (VTOCIX) or rebuild a disabled VTOCIX on a shared DASD volume *while it is online to multiple systems*. The SWAPBUILDIX function of FDRPAS uses FDRPAS communication and co-ordination techniques to allow the VTOCIX to be built while online to all sharing systems.

Summary

Storage Management is about more than just taking backups. The powerful utilities described in this section of the guide complement the main backup and restore features outlined in PARTS ONE and TWO to provide today's Storage Administrators with a *total* Storage Management System.

PART FIVE

STORAGE MANAGEMENT REPORTING

Introduction

A good set of reporting tools is an essential element in any Storage Management System. As you've already seen from the sample job output on page 29 (FASTCPK) and page 31 (FDRREORG), each of the modules of the FDR Family has extensive built-in reporting to keep you fully informed on what's been done, when, where and how. This information is vital for allowing you to control your Storage Management System and to ensure that what you think is being done really *is* being done.

Of course, it is also important to keep track of all Storage-related issues, not just the activities of the Storage Management System itself. At any computer installation, for example, staff at a variety of levels need to be able to refer to accurate and timely information about the use of Storage resources. Accurate information is the only basis for sound decision making and future planning. Without this, it is impossible to determine how efficiently Storage space is being utilized, or predict how needs will grow. Accurate information is essential for solving existing problems and, better still, for preventing problems from arising in the first place.

With these requirements in mind, the FDR Storage Management System offers several separate reporting modules. In summary, they are:

- FDRQUERY—A reporting and analysis tool that provides simulated results of Volume Incremental Backups and Data Migration *without* first having to run these processes for real.
- FDRABRP—A set of fixed-format reports to produce information on numerous aspects of the ABR component of the FDR Storage Management System.
- FDREPORT—A free-format, full function DASD Management reporting tool which can draw information from a variety of sources (*e.g.* VTOC, VVDS, Catalog etc.) to produce customized reports and summaries on all DASD-related issues.
- SRS—An online (TSO/ISPF) reporting tool with all the power and flexibility of FDREPORT.

Let's take a closer look now at each of the above reporting modules, starting off with FDRQUERY…

FDRQUERY

FDRQUERY is a statistics query program that can scan some/all online DASD volumes to determine which ones would gain the most benefit from:

- Volume Incremental Backup—as described in PART TWO, pages 8 and 9
- Data Migration—as described in PART FOUR, page 23.

On the Volume Incremental Backup report, FDRQUERY scans the selected DASD volumes and compares the number of tracks that would be dumped by a full-volume dump program (*e.g.* FDR) with the estimated 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.

PART FIVE

STORAGE MANAGEMENT REPORTING

Here is an example of an FDRQUERY Incremental backup job, together with some sample output. In this particular run, we are looking at all online DASD volumes beginning with PROD.

```
//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 occupying 7359 tracks. Only 91 (3947 tracks) currently have the MVS Update Indicator switched on. An incremental backup would make a savings of 46% in the amount of tracks backed up compared to a full-volume backup. This would have a direct effect on the elapsed time of the backup and the amount of tape required, reducing both by a similar percentage. PROD02 shows even greater potential for savings. It contains 97 datasets, occupying 11223 tracks. Only 18 (1634 tracks) have the MVS Update Indicator switched on, yielding an 86% savings. In the Summary report, the overall savings for the two PROD volumes would be 70% compared to full-volume backups.

PART FIVE

STORAGE MANAGEMENT REPORTING

In this next example, FDRQUERY is doing a Data Migration simulation on all volumes belonging to the DFSMS Storage Group called 'TSO'. The details for one volume (TSO002) are shown, together with the Summary report for all 6 volumes in the Storage Group.

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

                        REPORT  ARCHIVE,STORGRP=TSO
```

FDRQUERY INCREMENTAL BACKUP SAVINGS

VOLSER	DEVTYPE	ALLOC TRACKS	BEFORE %ALLOC	AFTER %ALLOC	LAST USED DAYS	LAST USED DATE	SAVINGS IF MIGRATED DSNS	SAVINGS IF MIGRATED TRACKS	SAVINGS IF MIGRATED %SAVED
TSO002	3380-K	36045	90.51%	43.19%	30	89051	6901	18841	52.27%
				54.94%	60	89021	4787	14162	39.29%
				63.71%	90	88356	2962	10672	29.61%
				68.58%	120	88326	1897	8730	24.22%

SUMMARY LEVEL BY DEVICE TYPE

VOLCNT	DEVTYPE	ALLOC TRACKS	BEFORE %ALLOC	AFTER %ALLOC	LAST USED DAYS	LAST USED DATE	SAVINGS IF MIGRATED DSNS	SAVINGS IF MIGRATED TRACKS	SAVINGS IF MIGRATED %SAVED
6	3380-K	173095	72.44%	52.66%	30	89051	9423	47264	27.30%
				55.81%	60	89021	5296	39722	22.94%
				60.48%	90	88356	4199	28558	16.49%
				64.94%	120	88326	2972	17911	10.34%

TSO002 is currently 90% allocated. If Data Migration was run against the volume to remove all datasets not referenced in the last 30 days, the allocation would come down to 43%. 6901 datasets would be removed from the volume (18,841 tracks), saving 52% of the allocated space. If a less severe criterion was used for the migration, the savings would not be as high. The report shows the results of using 'not-referenced-in' criteria of 60, 90 and 120 days. The Summary report shows that, across all 6 volumes in the TSO Storage Group, the current allocation level is 72%. With 30-day migration criterion, this would be reduced to 52%, with over 9,000 datasets (nearly 50,000) tracks being reclaimed.

PART FIVE

STORAGE MANAGEMENT REPORTING

FDRABRP

The FDRABRP reporting utility provides ABR users with a set of simple, fixed-format reports. Although, in general, we recommend the use of the SRS panels for most reporting needs (see “S.R.S.” later), these batch-driven FDRABR reports can occasionally be useful.

In this first example, we are reporting on all current backups for TEST.DSN on volume PROD01. As you can see, there are several backups currently being recorded by ABR. The report shows the date of each backup, the unique suffix of the backup file, and also the tape volser(s) on which the backup is stored.

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

     PRINT   BACKUP,DSN=TEST.DSN
```

```
                    FDRABR DATASET BACKUP REPORT

             VOLUME  VL D/S   TRK          BACKUP  BK       TAPE   TAPE
DSN          SERIAL  SQ ORG  ALLOC FREE      DATE  NO     SUFFIX   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 this second example (using similar JCL but a different PRINT command), FDRABRP has reported on all migrated datasets belonging to the ‘USTEST.ZZ75’ application. For each migrated dataset, various pieces of information are shown, including the date of the migration, the volser of the DASD volume on which the dataset was originally stored, the expiration dates of both copies of the migrated dataset (one on disk and one on tape) and also the suffix and tape volser(s) where the migrated data now resides.

```
                                FDRABR MIGRATION REPORT

*-DSK VOL-* *--DATE OF--* D/S REC BLOCK LBP BK **---DSN TAPE INFO---**
DATASET NAME              SERIAL   S  TY   ARCHIVE  EXPDT      ORG  FM    SIZE  LRECL ALLOC  TY    CN    SUFFIX  FN VOLSER(S)
------- ----              ------   -  --   -------  -----      ---  --    ----  ----- -----  --    --    ------  -- ---------

USTEST.ZZ75.IDPSEQ.LIBDAT IDPLB4   1  0E   2001.341 2002.006   PS   VB   32760  32756   1T   0E    1    B101341A   0  IDPBK0
                                                                                             80    2    B201341A   7  BA0077
USTEST.ZZ75.IDPSEQ.BIN    IDPLB4   1  0E   2001.341 2002.006   PS   FB   12004  12000 300T   0E    1    B101341A   0  IDPBK0
                                                                                             80    2    B201341A   7  BA0077
USTEST.ZZ75.IDPSEQ.CHNGE  IDPLB2   1  0E   2001.334 2002.364   PO   FB   19040     80   9C   0E    1    B101334A   0  IDPBK0
                                                                                             80    2    B201334A  15  BA0076
USTEST.ZZ75.IDPSEQ.DATS   IDPLB0   1  0E   2001.334 2002.364   PO   FB    9040     80  89T   0E    1    B101334A   0  IDPBK0
                                          2003.365                                           80    2    B201334A  15  BA0076
```

PART FIVE

STORAGE MANAGEMENT REPORTING

The FDRABRP reports can either be produced by running a simple batch job (as the examples above), or they can be requested through the ABR ISPF panels:

```
---- FDR TOTAL DASD MANAGEMENT SYSTEM - FDRABR REPORT PANEL ----

REPORT OPTION ===>

BLANK  -  ARCHIVE            ENTER 'C' TO CHANGE FORMAT ===>
  2    -  BACKUP             PRINT ALL AVAILABLE BACKUPS
  3    -  SCRATCH            CREATE DETAIL AND SUMMARY REPORTS
  4    -  CATALOG            REPORT DIRECT TO TSO
  5    -  VOLUME STATUS      DISPLAY IN TSO FORMAT
  6    -  FDREPORT           CREATE 58 LINES PER PAGE

FDREPORT NAME ===>

PROJECT ===>
LIBRARY ===>
QUALIFIER =>

OTHER DATASET NAME ===>
OTHER DATASET GROUP ==>

VOLUME SERIAL ===>
```

FDREPORT

FDRQUERY and FDRABRP are extremely useful reporting tools, but they are both designed to produce fixed-format reports with very specific uses. FDREPORT, on the other hand, is a powerful, free-format DASD Management Reporting tool that can quickly and efficiently generate reports on a wide range of DASD related issues, either through the pre-defined “Health Check” reports (see later), or by running customized, site-specific reports.

FDREPORT can access information from a variety of sources, including VTOCs, VVDS’s, Catalogs, PDS Directories, ABR Control Files and the DFHSM MCDS and BCDS.

It also generates additional information (such as %Free figures) which may not be directly available, but are derived or calculated from other sources.

PART FIVE

STORAGE MANAGEMENT REPORTING

Output can be produced in various formats, including screen-based reports, printed reports, and even 'punched' machine-readable data or job streams.

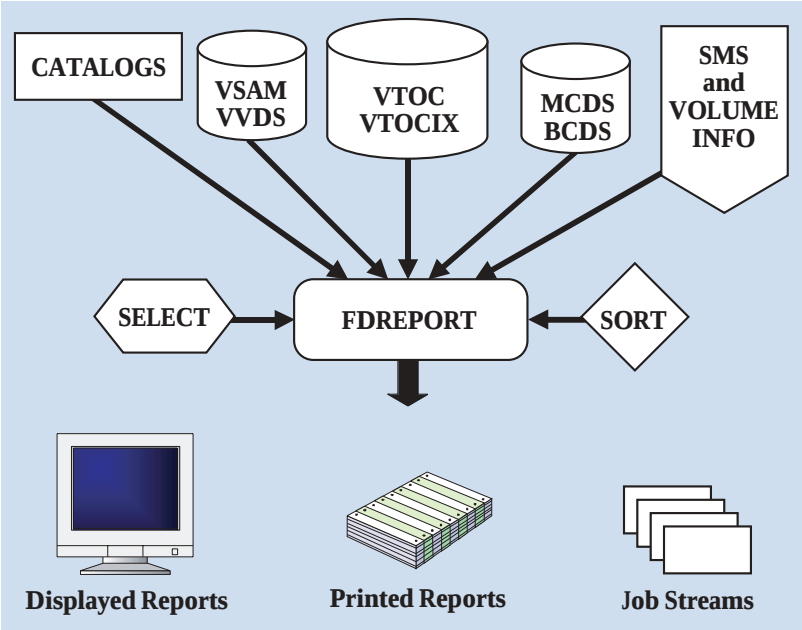

FDREPORT can be used to produce very detailed output, (*e.g.* for a Storage Administrator), or it can produce concise, summarized reports destined perhaps for Capacity Analysts or High-level Management.

This report identifies datasets likely to suffer Sx37 space abends.

It shows all datasets with less than 10% freespace (with no secondary allocation), all non-VSAM datasets with 13 or more extents, and all VSAM datasets with 100 or more extents.

FDREPORT—Example #1

```
//REPX37   EXEC PGM=FDREPORT,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//ABRMAP    DD  SYSOUT=*
//SYSIN     DD  *

    TITLE     LINE='DATASETS THAT MAY RUN OUT OF SPACE'
    XSELECT   SECALLOC.EQ.0,%FREE.LT.10
    XSELECT   DSORG.NE.EF,NOEXTENT.GE.13,%FREE.LT.10
    XSELECT   DSORG.EQ.EF,NOEXTENT.GE.100,%FREE.LT.10
    REPORT    FIELD=(DSN,DSORG,SECALLOC,%FREE,NOEXTENT,SIZE)
    PRINT     ENABLE=ONLINE
```

```
                    DATASETS THAT MAY RUN OUT OF SPACE

DSN               DSORG   VOLSER   SECAL   %FREE   EXTENTS   ALLOC
---               -----   ------   -----   -----   -------   -----
MASTER.FILE       EF      PROD12   10          7       105     780
PROD.CNTL.CARDS   PO      SYS123   0           3         1      20
```

PART FIVE

STORAGE MANAGEMENT REPORTING

This next example shows FDREPORT's ability to access and report on information in the DFSMShsm MCDS and BCDS databases.

The report shows migrated datasets, selected via the high-level qualifier of 'PLOUS'.

FDREPORT—Example #2

```
//REPHSM   EXEC PGM=FDREPORT,REGION=0M
//SYSPRINT  DD  SYSOUT=*
//ABRMAP    DD  SYSOUT=*
//MCDSDD    DD  DSN=HSM.MCDS,DISP=SHR
//SYSIN     DD  *

    XSELECT   XDSN=PLOUS**
    REPORT    FIELD=(DSN,VOL,TVTOCDSN)
    PRINT     DATATYPE=MCDS
```

```
DSN               VOLSER   TVTOC BACKUP FILE NAME
---               ------   ---------------------
PLOUS.PS.AV001    RVA7C8   HSM.HMIG.T370611.PLOUS.J0256
PLOUS.PS.AV801    RVA7C8   HSM.HMIG.T420611.PLOUS.J0256
PLOUS.RVA7C8.DOC  RVA7C8   HSM.HMIG.T440611.PLOUS.RVA7C8.J0256
```

As well as being able to produce an almost infinite range of free-format reports, FDREPORT includes a set of pre-defined reports. Their purpose is two-fold:

- **To provide real-life examples of the facilities available within FDREPORT**

  As well as the standard features, like Selection, Sorting and Summarizing, the pre-defined reports also make extensive use of some of the more advanced FDREPORT features.

- **To provide a HEALTH CHECK on your installed DASD**

  All of the pre-defined reports have been designed to look for specific problems, both at the *dataset level* and the *volume level*. When run as a whole, they can provide a complete Health Check of your installed DASD.

The reports (numbered HCHECK1—HCHECK8) can be run with a minimum of change—usually only requiring the addition of an appropriate job card. All of the reports default to reporting against all ONLINE DASD volumes. However, because they are constructed using standard FDREPORT control statements and JCL, they can easily be tailored to provide more specific or targeted reporting, if required.

PART FIVE

STORAGE MANAGEMENT REPORTING

The tables below list of the type of problems highlighted by each of the Health Check reports:

HCHECK1—DASD With Potential Problems:

- Disk Volumes More Than 80% Full
- IBM Fragmentation Index, Worst First
- Volume Mount & Use Status
  SMS, VTOCIX Status
- VTOC/VVDS/VTOCIX's More Than 80% Full
- VVDS In Multiple Extents
- VTOCs/VVDS's with Logical Errors

HCHECK2—Reports For SMS Administrator

- SMS Volume status
- SMS Volumes with disabled VTOCIX's
- Dataset's SMS Attributes
- Information On PDSE Datasets
- Uncataloged Datasets On SMS Volumes
- Non-Managed Datasets On SMS Volumes
- Datasets Ineligible For SMS Management

HCHECK3—Reports For Performance Analyst

- Disks With Potential Performance Problems
- Multi-Extent VSAM Datasets
- Multi-Extent Non-VSAM Datasets
- Multi-Volume Datasets
- Datasets Likely To Get Sx37 Abends

HCHECK4—Reports For Capacity Analyst

- Disks With Potential Wasted Space
- Total Space Available, allocated and used
- Overallocation In All Datasets
- Datasets Not Referenced in 60 Days
- Datasets With Inefficient Blocksizes

HCHECK5—Reports For VSAM Tuning Analyst

- Space Occupied By Largest VSAM Files
- Clusters With High Inserts
- Clusters with high CI/CA splits
- Clusters With 3 Or More Index Levels
- Multi-Extent, Multi-Volume VSAM Files
- Overallocated Clusters

HCHECK6—Reports On PDS Efficiency

- Basic Stats On PDS/PDSE Datasets
- Multi-Extent PDS/PDSE Datasets
- PDS's With Limited Freespace
- PDS's With Limited Free Directory Blocks
- PDS's With Excess Free Space
- Space Reclaimable By PDS Compression

HCHECK7—Datasets That Offend Standards

- Uncataloged Datasets
- Wrongly cataloged datasets
- Datasets With Undefined DSORG
- Empty PS, PO, VSAM Datasets
- Unmoveable Datasets
- Datasets existing as only a catalog entry

HCHECK8—Project Or DASD Mgmt Reports

- Summary Of Dataset Types
- Space Occupied By An Application
- Space Summarized By Index
- DASD Utilization By Esoteric Unit Name
- DASD Utilization By SMS Storage Group
- DASD Utilization By Volser Group, devtype.

PART FIVE

STORAGE MANAGEMENT REPORTING

S.R.S

In addition to the batch-driven reporting illustrated earlier, the power and flexibility of FDREPORT is also available under TSO via the ***Search, Reporting and Services*** dialog, or ***SRS*** for short. The SRS panels permit a 'menu-driven' execution of FDREPORT, with the same wide range of selection criteria as a batch-driven FDREPORT. Data can be extracted from the same sources, with full control over the content, sort order and layout. The selection criteria can be saved for later re-use. In the example shown here, the user has requested a report showing the Creation Date of all datasets beginning with RTS which currently reside on 3390-3 DASD.

```
       -------  S R S DATASET SELECTION  -------
COMMAND ==>                            SCROLL ==> PAGE

FIELD            SELECTION VALUE          REPORT   SORT
SOURCE/FIELDS
DSNAME     ==>   RTS**                      1
VOL        ==>                              3
SOURCE     ==>   CATALOG (Catalog Volume Arc..)
CATALOGN   ==>
ARCDSN     ==>
VTOC FIELDS
DEFAULTS   ==>
UNIT       ==>
DEVTYPE    ==>   3390-3
VOLSQ      ==>
DATES      ==>
CRDATE     ==>                              2        1A
```

The resultant report will contain three columns—Dsname, Crdate and Vol—and the data will be displayed in ascending order of Creation Date.

```
       -------  S R S DATASET LIST PANEL  -------
COMMAND ==>                            SCROLL ==> PAGE

COMMAND    DATASET NAME             CRDATE          VOLSER

           RTSAL.FDREPORT.CNTL      1998.036        MVS001
           RTS.JOB.CNTL             1998.234        TSO004
           RTS.OLD.FILE             1999.010        DEV020
           RTS.VSAM.FIL             1999.156        MIGRAT
           RTS.DUMP.DSN             1999.254        DEV017
           RTS1.OTHER.FILE          1999.315        TSO010
           RTS.PDS                  2000.003        DEV012
           RTS.PDSE.DSN             2000.121        MIGRAT
```

On the far left-hand side of the above panel, you will see that there is a COMMAND field, just like a normal ISPF display. This command field can be used to invoke a variety of functions for one or more of the datasets that are being displayed on the dataset list. These commands include:

- SRS services, such as 'I' (Info), 'M' (Member List)
- FDR/ABR services, such as REORG, RECALL, COPY
- Normal ISPF functions, like Edit and Browse
- TSO commands, including DELETE, LISTDS, LISTCAT and RENAME
- CLISTs and REXX execs

These commands can all be used in the left-hand 'command' column.

PART FIVE

STORAGE MANAGEMENT REPORTING

It is also possible to produce ‘volume-level’ reports with SRS. The following example reports on all volumes in the SMS Storage Group called ‘General’ that have a volser starting ‘PROD’. The report will show the disk volser, the percentage of free space on the volume and the percentage of free space within the VTOC. The output will be sorted so that volumes with the most freespace appear first.

```
        -------  S R S VOLUME SELECTION  -------
COMMAND ==>                          SCROLL ==> PAGE

FIELD           SELECTION VALUE          REPORT   SORT

VLVOLSER  ==>   PROD*____________________   __1__   ____
VLUNIT    ==>   ____________________________ ____   ____
VLDEVTYP  ==>   ____________________________ ____   ____
STORGRP   ==>   GENERAL__________________    ____   ____

VL%FTRKS  ==>   ____________________________ __2__   _1D_
VL%UTRKS  ==>   ____________________________ ____   ____
VL%FDSCB  ==>   ____________________________ __3__   ____
VL%UDSCB  ==>   ____________________________ ____   ____
VL%FINDX  ==>   ____________________________ ____   ____
VL%UINDX  ==>   ____________________________ ____   ____
VLBYTRK   ==>   ____________________________ ____   ____
VLALOTRK  ==>   ____________________________ ____   ____
```

The resulting report will look like the one below. Clearly, the three volumes at the top of the report are very under-allocated and represent a potential waste of disk space. Another volume (PROD21) is only about half-full, but it has very limited free space remaining in the VTOC. It will soon be impossible to allocate new datasets on this volume, even though there are ample free tracks.

```
        -------  S R S VOLUME SELECTION  -------
COMMAND ==>                          SCROLL ==> PAGE

COMMAND       VOLSER          %FTRKS          %FDSCB

              PROD42          95              80
              PROD83          91              67
              PROD03          89              54
              PROD82          56              36
              PROD21          46              03
              PROD45          31              25
              PROD03          25              76
              PROD05          20              11
```

Summary

The superior performance of reporting tools like FDREPORT and SRS allow the data center administrators and the end-users to quickly obtain (in real time) the information they require to make their day-to-day decisions regarding Storage Management. These reports can be generated in a matter of minutes, even when collecting information from several hundred DASD volumes.

PART FIVE

STORAGE MANAGEMENT REPORTING

SYSTEM SUMMARY

The following diagram summarizes the various modules within the FDR Storage Management Family.

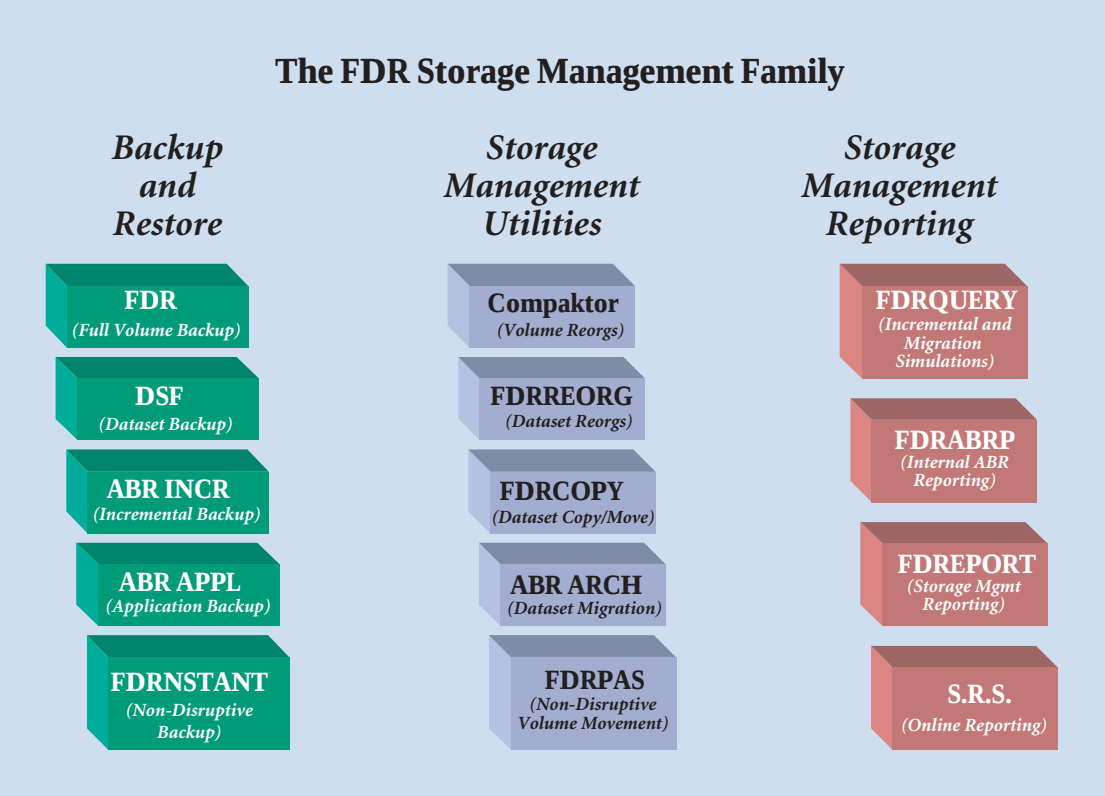

Many of the components of the FDR Storage Management Family can be licensed individually, thereby allowing Storage Management issues to be addressed efficiently and cost-effectively. Contact your local Innovation office or representative for further information.

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