




Storage Subsystem Library

GC26-4492-1

**Using the IBM 3380
Direct Access Storage
in an MVS Environment**





**Using the IBM 3380
Direct Access Storage
in an MVS Environment**

Federal Communications Commission (FCC) Statement

Warning: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instructions manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

Instructions to User: Properly shielded and grounded cables and connectors must be used for connection to peripherals in order to meet FCC emission limits. Proper cables are available from IBM authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables or by unauthorized modifications to this equipment. It is the responsibility of the user to correct such interference.

Second Edition (October 1988)

This edition replaces and makes obsolete the previous edition, GC26-4492-0.

The changes for this edition are summarized under "Summary of Changes" following the preface. Specific changes are indicated by a vertical bar to the left of the change. These bars will be deleted at any republication of the page affected. Editorial changes that have no technical significance are not noted.

Changes are made periodically to this publication; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370, 30xx, and 4300 Processors Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM licensed program in this publication is not intended to state or imply that only IBM's program may be used. Any functionally equivalent program may be used instead.

Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality. If you request publications from the address given below, your order will be delayed because publications are not stocked there.

A readers' comment form is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Programming Publishing, P.O. Box 49023, San Jose, California, U.S.A. 95161-9023. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Trademarks

The following names have been adopted by IBM for trademark use and will be used throughout this publication:

ESA/370™

Enterprise System Architecture/370™

MVS™

MVS/DFP™

MVS/ESA™

MVS/SP™

MVS/XA™

Processor Resource/Systems Manager™

PR/SM™

VM™

3090™

Contents

Chapter 1. Introducing the IBM 3380 Direct Access Storage	1
3380 Direct Access Storage Overview	1
Attaching the 3380 to Storage Controls and Processors	2
Chapter 2. Planning for Installation	5
Physical Planning	5
Identifying 3380 Physical Requirements	5
Reserving I/O Addresses	5
Adding 3380 Strings to the Operating System	6
Identifying Prerequisite Features	6
Software Planning	7
Planning for Programming Support	7
Planning for Software Tool Availability	11
Plan to Apply Software Fixes	15
Planning for Publications Availability	16
Administrative Planning	16
Creating a Project Team and Project Plan	17
Verifying Device Delivery Schedules	17
Scheduling System Time for Volume Preparation and Data Movement	17
Training Personnel	17
Documenting New Procedures	18
Chapter 3. Understanding Your Current Configuration	19
Determining Data Set Inventory	19
MVS Tools for Determining Data Set Existence, Placement, and Characteristics	20
Identifying Performance-Critical Data Sets	21
Frequently Accessed System Data Sets	21
Frequently Accessed DFHSM Data Sets	21
Frequently Accessed IMS/VS Data Sets	21
Frequently Accessed DB2 Data Sets	22
Frequently Accessed User Data Sets	22
Measuring the Effectiveness of Your Current Hardware Configuration	22
Using RMF to Review I/O Workload	22
RMF Measurements	23
Sorting Current Volumes by I/O Activity	23
Removing Device-Dependent Code	24
Where to Look for Device-Dependent Code	25
Identifying Device Dependencies	26
Techniques to Minimize Device-Dependent Code	26
Chapter 4. Planning the Hardware Configuration	29
Understanding Hardware Capabilities	29
Internal Paths	29
Dynamic Path Selection (DPS)	31
Device Level Selection (DLS)	32
Device Level Selection Enhanced (DLSE)	32
Speed Matching Buffer for 3380	33
Access Arm Movement	33
Channel Switches	34
Capacity Considerations	34
Performance Considerations	35

I/O Activity	36
Availability Considerations	37
Shared DASD Considerations	38
Guest System Considerations	38
Storage Control Selection	39
Cache Storage Controls	39
Configuring 3380 Strings	40
Planning the Hardware Configuration	41
Basic 3380-3880 Configuration	42
Intermixed Strings on 3990 Single-path Storage Directors	43
Four-Path and Two-Path String Intermix	44
A 3380-3990 4-Path Configuration for High Availability	45
CJ2 Configuration	46
Documenting the Hardware Configuration	46
Chapter 5. Planning the Data Configuration	47
Identifying Data to Move or Create	47
Volume Table of Contents (VTOC)	47
Work Files	48
Page and Swap Data Sets	48
JES2 Spool Data Sets	48
JES2 Checkpoint Data Sets — MVS/SP 2.1.5	49
JES2 Checkpoint Data Sets — MVS/SP 2.2.0	49
JES3 Spool Data Sets	50
JES3 Data Sets	51
Other System Data Sets	51
Catalogs	52
IMS/VS Data Sets	53
DB2 Data sets	53
VSAM Application Data Sets	53
Non-VSAM Application Data Sets	53
Unmovable Data Sets	53
Data Sets Created by Non-IBM Programs	54
Tape Data Sets	54
Considerations for the Placement of Data Sets	54
Planning a Data Management Strategy	55
Using Storage Pools	56
Using Storage Groups	56
Documenting Configuration and Data Set Placement	57
Space Utilization Guidelines	57
Determining Space Requirements	58
Developing a Data Movement Strategy	60
Planning the Data Movement Sequence	60
Alternate Methods to Put Data on 3380s	62
Planning for Backup and Recovery While Moving Data	63
Documenting and Reviewing Your Data Movement Strategy	63
Scheduling System Time	64
Chapter 6. Installing the 3380	65
Defining the 3380 to the System	65
Assigning I/O Addresses or Device Numbers	65
Defining the I/O Configuration to MVS	66
Installing the 3380 Units	73
Preparing the Volumes for Use	73
Calculating the Size of VTOC and VTOC Index	73
Initializing New Volumes	74

Chapter 7. Operating the 3380 under MVS	77
Using the DISPLAY Command to Determine Device Status	77
The DISPLAY U Command	78
The DISPLAY M Command in MVS/XA and MVS/ESA	79
The DISPLAY M Command in MVS/370	79
The DISPLAY Command with ALLOC Parameter	80
Using the DEVSERV Command in MVS/XA and MVS/ESA	81
Using the VARY Command	81
The VARY Command for Devices	83
The VARY PATH Command	83
The VARY PATH Command and Path Group ID	84
Missing Interrupts	85
Chapter 8. Moving Data in the MVS Environment	87
Modifying Examples for Your Use	87
Backup and Recovery While Moving Data	88
Creating a Backup Tape	88
Restoring a Volume	88
Restoring Selected Data Sets	88
Using the DFDSS COPY Command	89
Data Set COPY	89
Full-Volume COPY	95
VTOC Considerations	95
One 3380 Volume to Another 3380 Volume	96
Three Standard Capacity 3380 Volumes to One Triple Capacity 3380 Volume	97
Three Double Capacity 3380 Volumes to Two Triple Capacity 3380 Volumes	98
Two 3350 Volumes to a 3380 Volume	99
Four 3350 Volumes to a Double Capacity 3380 Volume	100
Six 3350 Volumes to a Triple Capacity 3380 Volume	101
Using DUMP/RESTORE to Move Data	101
Moving Data Sets to Smaller Capacity 3380 Models	102
Moving Tape Data Sets	102
Chapter 9. Moving IMS/VS and DB2 Data Sets	103
Moving IMS/VS Data	103
Moving DB2 Data	106
Considerations Before Moving DB2 Data Sets	106
Moving a DB2 Data Set or an Entire DASD Volume to Another Volume	106
Considerations Before Moving Storage Groups	108
Procedures for Moving DB2 Table Spaces and Indexes	108
Chapter 10. Performing Follow-Up Activities	111
Tuning the System for Performance	111
Backing Up Your Data	112
Data Base Backup	112
Incremental Backup	113
Backup and Recovery Considerations	113
Full-Volume Backup and Recovery	113
Documenting the New Configuration and Data Set Placement	113
Updating Written Operating Procedures	114
Performing Media Maintenance	114
System Exception Reports	115
Using the Service Information Messages	115
Interpreting the SIM Alert Message	115
Establishing SIM Handling Procedures	116
Generating EREP Report Requests	116

Appendix A. Sample DASD Installation and Migration Project Plan	121
Acronym List	131
Glossary	133
Bibliography	145
Index	151

Figures

1.	The Storage Subsystem Library	xv
2.	Storage Subsystem Library Restructure	xx
3.	IBM Processors and Storage Controls Compatible with 3380 Strings	3
4.	Minimum Programming Support in the MVS/XA Environment	8
5.	Minimum Programming Support in the MVS/ESA Environment	9
6.	Minimum Programming Support in the MVS/370 Environment	10
7.	Software Tool Support in the MVS Environment	12
8.	3380 AA4 Internal Data Path Sharing	30
9.	Summary of Data Capacity for 3380 Models	35
10.	Storage Control Cache Sizes	39
11.	Example of 3380 Model AE4 and AK4 Strings Attached to a 3880	42
12.	Intermixed 3380 2-path Strings Attached to a 3990 Model 2 or 3	43
13.	3380 4-Path and 2-Path Strings Intermixed on a 3990 Model 2 or 3	44
14.	3380 Model AJ4 and AK4 4-Path Strings Attached to a 3990 Model 2 or 3	45
15.	Example of 3380 Model CJ2 String Attached Directly to Channel	46
16.	Using Average Record Length and the AVGREC Subparameter	59
17.	Configuration Assumed for MVS System Generation Examples	67
18.	Sample 3380 IODEVICE Statement	68
19.	IOCP Generation, MVS/370 Operating System with 370 Extended Architecture Processor	68
20.	IODEVICE Statement for MVS/XA or MVS/ESA	69
21.	IOCP Generation, MVS/XA or MVS/ESA	69
22.	Four-Path Configuration Example	70
23.	IOCP Generation for a Four-Path 3380-3990 Configuration	70
24.	IOCP Generation, 3380 CJ2 String	71
25.	Four-Path, Two-Path Intermixed String Configuration Example	72
26.	IOCP Generation for Four-Path, Two-Path Intermixed String	72
27.	Offline Initialization of 3380 Volume	75
28.	Example of DISPLAY U Command	78
29.	Example of DISPLAY M=CHP Command in the MVS/XA and MVS/ESA Environments	79
30.	Example of DISPLAY M=DEV Command in the MVS/XA and MVS/ESA Environments	79
31.	Example of DISPLAY M Command in the MVS/370 Environment	80
32.	Example of D u,,alloc Command	80
33.	Example of DEVSERV Command in the MVS/XA and MVS/ESA Environments	81
34.	Using DFDSS to Move an Integrated Catalog Facility User Catalog	90
35.	Using DFDSS to Move Selected Unmovable Data Sets	91
36.	Using DFDSS to Move a Multivolume Data Set	92
37.	Using DFDSS to Move System Data Sets	94
38.	Using DFDSS to Copy by Volume from 3380 to 3380	96
39.	Using DFDSS to Move Three Standard Capacity 3380 Volumes to a Triple Capacity 3380 Volume	97
40.	Using DFDSS to Move Three Double Capacity 3380 Volumes to Two Triple Capacity 3380 Volumes	99
41.	Using DFDSS to Move Two 3350 Volumes to a 3380 Volume	100
42.	Using DFDSS to Move Four 3350 Volumes to a Double Capacity 3380 Volume	100
43.	Using DFDSS to Move Six 3350 Volumes to a Triple Capacity 3380 Volume	101
44.	Moving a Tape Data Set to 3380 Direct Access Storage	102

45.	MVS/370, MVS/XA, and MVS/ESA SIM Alert Message Format	115
46.	Meaning of SIM Alert Severity Field for SCU Failing Component	116
47.	Sample System Exception Report Request	117
48.	Example of a Formatted Detail Edit and Summary Report	117
49.	Example of an Unformatted Detail Edit and Summary Report	118
50.	Example of SIMs in System Exception Informational Messages	119
51.	Example of 3380 Model AE4 and AK4 Two-Path Strings Attached to a 3880	141
52.	Example of 3380 Two-Path Strings Attached to a 3990 Model 2 or 3 in DLS Support Mode	142
53.	Example of 3380 Enhanced Subsystem Model Four-Path Strings Attached to a 3990 in DLSE Support Mode	143
54.	Example of 3380 Two-Path and Four-Path Strings Attached to a 3990 in DLSE Support Mode	144

Preface

This manual is part of the Storage Subsystem Library (SSL)—a set of manuals that provides information about the hardware components of IBM disk storage subsystems. The SSL includes both direct access storage devices (DASD) and storage control publications, this manual is part of the SSL subset that is concerned primarily with 3380 DASD.

To use this manual effectively, you should also plan to read the SSL manuals, *IBM 3380 Direct Access Storage Introduction*, for a description of all 3380 models except the 3380 Model CJ2, *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, for a description of the 3380 Model CJ2, and *Maintaining IBM Storage Subsystem Media*, for information on 3380 maintenance. The manuals describing the storage controls to which you will attach the 3380 should be available for reference. You should also have a complete set of publications for your MVS operating system available for reference.

About This Manual

This manual is written for the MVS systems programmer, storage administrator, or hardware specialist responsible for installing and using the 3380 in the MVS/ESA, MVS/XA and MVS/370 environments.

In addition to device-specific information for the various models of the 3380, this manual also illustrates techniques for more efficient storage management under MVS. It offers guidance on managing system performance, availability, and space through effective use of the DASD storage subsystem.

This manual contains:

Chapter 1, “Introducing the IBM 3380 Direct Access Storage” on page 1, provides a brief description of IBM 3380 model groups and strings. Attachment options for each 3380 model to storage controls and processors are also presented.

Chapter 2, “Planning for Installation” on page 5, lists the activities involved in the physical installation of 3380s. This chapter helps you identify physical, software, and administrative areas that require planning prior to installation.

Chapter 3, “Understanding Your Current Configuration” on page 19, helps you determine the types of data you have on your system. It describes tools you can use to inventory data sets and to measure DASD utilization. Types of MVS data to move and guidelines on when to move data sets are presented.

Chapter 4, “Planning the Hardware Configuration” on page 29, helps you determine how to effectively use new volumes. This chapter presents hardware, capacity, and performance related items to consider when planning your new hardware configuration. Configuration diagrams are included to show you how to configure 3380s for performance and availability.

Chapter 5, “Planning the Data Configuration” on page 47, provides information that will help you decide where to place data sets on your 3380 devices. The implementation of **system-managed storage** using storage pools and storage groups is discussed. This chapter describes a sequence for moving data and the need for a backup and recovery strategy.

Chapter 6, "Installing the 3380" on page 65, helps you install the 3380 on your system and prepare the volumes for use. System generation, hardware installation, VTOC calculation, and device initialization are described.

Chapter 7, "Operating the 3380 under MVS" on page 77, shows you how to use MVS system commands to determine and change the status of your 3380 devices.

Chapter 8, "Moving Data in the MVS Environment" on page 87, shows you how to use DFDSS to move data onto your 3380. DFDSS background information is provided as well as specific examples of using DFDSS commands to move data sets from one volume to another.

Chapter 9, "Moving IMS/VS and DB2 Data Sets" on page 103, provides guidelines and procedures to use when moving IMS/VS and DB2 data.

Chapter 10, "Performing Follow-Up Activities" on page 111, describes follow-up activities performed after device installation. These activities include documenting system operating procedures and data management tasks. Maintaining the DASD media and documenting your new configuration is also discussed.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121, provides a sample device installation and migration project plan that you can use as a model for your data processing center.

"Acronym List" on page 131 contains definitions for acronyms.

"Glossary" on page 133 lists the terms and abbreviations used in the Storage Subsystem Library manuals.

"Bibliography" on page 145 lists the related publications that you can reference for further information on related topics and hardware. A brief description of their contents is provided to help you select the publications that are best suited to your needs.

Terminology

A comprehensive glossary is provided at the back of this manual. This glossary contains terms used not only in this manual but also terms, abbreviations, and acronyms from other DASD and storage control manuals in the Storage Subsystem Library.

Before reading further, be sure you understand how the following terms are used within this manual:

3380 refers to all models of the IBM 3380 Direct Access Storage, unless otherwise indicated.

Controller refers to the part of the 3380 A-unit or C-unit that controls the transfer of data between the devices and the storage control hardware.

Device refers to a uniquely addressable part of the 3380 unit that includes a set of access arms, their associated disks, and the electronic circuitry needed to locate, read, and write data.

MVS refers to all MVS/ESA, MVS/XA and MVS/370 operating systems. Unless specifically mentioned, the information in this book does not apply to OS/VS1. For more information on MVS/ESA see Chapter 2, "Planning for Installation" on page 5.

Storage Control refers to the hardware that handles interactions between the processor channel and the controller.

Volume refers to the set of disk surfaces associated with a device.

The Storage Subsystem Library

The Storage Subsystem Library describes characteristics, capabilities, and features of the hardware and provides instructions for installing, using, and maintaining storage subsystem components effectively in the various operating environments. Subsets of the library provide hardware- and software-related information for both DASD and storage controls.

SSL DASD Publications

The DASD subset of the Storage Subsystem Library contains the following manuals:

1. *IBM 3380 Direct Access Storage Introduction*, GC26-4491

Provides a complete description of the various models of the 3380, including characteristics, features, and capabilities. In addition, the configuration and attachment options are described along with other information that helps in designing a storage subsystem to meet your needs. The manual does *not* cover the 3380 Model CJ2.

2. *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, GC26-4497

Provides a complete description of the 3380 Direct Channel Attach Model CJ2 characteristics, features, capabilities, and string configuration options.

Provides specific guidance for using the 3380 in an MVS/ESA, MVS/XA or MVS/370 operating environment. The manual provides detailed instruction for planning the addition of new 3380 devices from a logical and physical point of view, installing devices, moving data to new devices, and performing ongoing activities to maintain a reliable storage subsystem.

4. *Using the IBM 3380 Direct Access Storage in a VM Environment*, GC26-4493

Provides specific guidance for using the 3380 in a VM/SP, VM/SP HPO, or VM/XA SF operating environment. The manual provides detailed instruction for planning the addition of new 3380 devices, installing devices, moving data to new devices, and performing ongoing storage management activities to maintain reliable performance and availability. In addition, storage considerations related to guest systems are addressed.

5. *Using the IBM 3380 Direct Access Storage in a VSE Environment*, GC26-4494

Provides specific guidance for using the 3380 in a VSE operating environment. The manual provides instruction for planning the addition of new 3380 devices, installing devices, moving data to new devices, and performing ongoing storage subsystem management.

6. *Maintaining IBM Storage Subsystem Media*, GC26-4495

Describes how the storage subsystem and the various operating systems handle disk storage errors and provides instruction on using the Environmental Record Editing and Printing (EREP) program and the Device Support Facilities (ICKDSF) program to diagnose and correct disk media errors. Recovery procedures are provided for the various device types. In addition, background material on DASD storage concepts is included.

7. *IBM 3380 Direct Access Storage Reference Summary*, GX26-1678

Provides a summary of 3380 capacity, performance, and operating characteristics.

SSL Storage Control Publications

The Storage Control subset of the Storage Subsystem Library contains the following manuals:

1. *IBM 3990 Storage Control Introduction*, GA32-0098

Provides a complete description of the various models of the 3990 Storage Control, including its data availability, performance, and reliability improvements over previous storage controls. In addition, the manual provides descriptions of the configuration attachment options, optional features, performance characteristics, and software support of the 3990 Storage Control.

2. *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*, GA32-0100

Provides a functional description of the 3990 Storage Control. The manual describes the planning, program installation, and storage management tasks used in typical environments. Configuration examples as well as sample programs for controlling the various functions of the 3990 Storage Control are provided.

3. *IBM 3990 Storage Control Reference*, GA32-0099

Provides descriptions and reference information for the 3990 Storage Control. The manual contains channel commands, error recovery, and sense information.

4. *Cache Device Administration*, GC35-0101

Specifies the access method services tools for administering a cache device under MVS. The book supports storage controls: 3990 Model 3, 3880 Model 23, 3880 Model 21, 3880 Model 13, and 3880 Model 11.

The *Storage Subsystem Library Master Index*, GC26-4496, provides a central source for information related to storage subsystem topics. Manuals for IBM 3380 Direct Access Storage, 3380 Direct Channel Attach Model CJ2, and 3990 Storage Controls are indexed in this publication. An overview of the material in the Storage Subsystem Library is provided with this index.

Figure 1 on page xv shows the relationships among the Storage Subsystem Library manuals in terms of high-level tasks described in each manual.

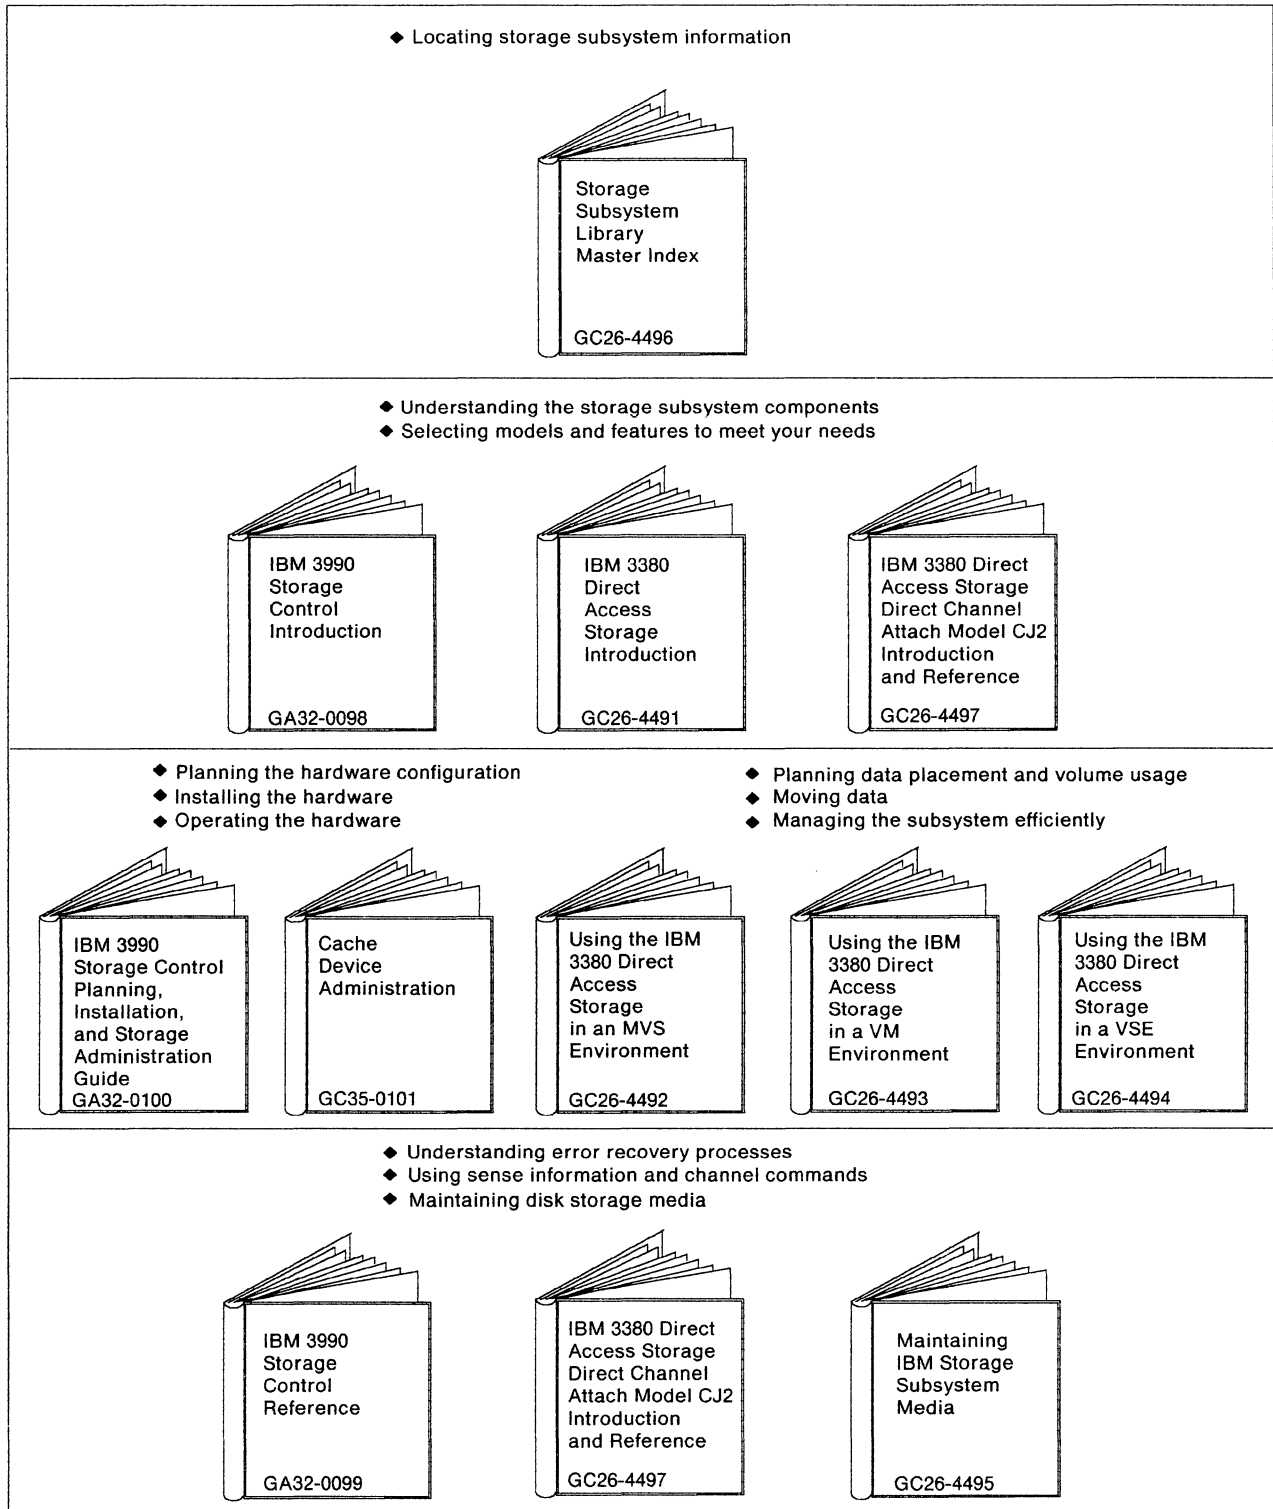


Figure 1. The Storage Subsystem Library

SSL Ordering Information

You can obtain a copy of **every manual** in the SSL using one General Bill of Forms (GBOF) number, **GBOF-1762**. Select one of the following GBOFs to obtain subsets of the SSL that provide information for specific environments and equipment. The GBOFs shown in the highlighted columns contain the manuals needed in the MVS environment. To obtain an individual SSL manual, use its order number.

Title	GBOF-1757	GBOF-1758	GBOF-1760	GBOF-1761	GBOF-0366
IBM 3380 Direct Access Storage Introduction, GC26-4491	X	X			
IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference, GC26-4497			X	X	
Using the IBM 3380 Direct Access Storage in an MVS Environment, GC26-4492					
Using the IBM 3380 Direct Access Storage in a VM Environment, GC26-4493	X		X		
Using the IBM 3380 Direct Access Storage in a VSE Environment, GC26-4494		X		X	
Maintaining IBM Storage Subsystem Media,* GC26-4495	X	X	X	X	
IBM 3380 Direct Access Storage Reference Summary, GX26-1678	X	X	X	X	
IBM 3990 Storage Control Introduction, GA32-0098					X
IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide, GA32-0100					X
IBM 3990 Storage Control: Reference, GA32-0099					X
Cache Device Administration, GC35-0101					X
Storage Subsystem Library Master Index, GC26-4496	X	X	X	X	X
Binder + 3380 inserts, GX26-3767	X	X	X	X	
Binder + 3990 inserts, GX26-3768					X

* Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage, GC26-4498, is distributed with this book.

SSL Binder Information

Binder kits for the SSL are available to help organize your library. Kits consist of a binder with identifying cover and spine inserts for either 3380 or 3990 subsets of the SSL, and are included when you order the following GBOF numbers:

- GBOF-1756 through -1761 each include a binder with 3380 inserts.
- GBOF-0366 includes a binder with 3990 inserts.
- GBOF-1762 includes binders and inserts for both 3990 **and** 3380.

SSL binder kits may also be ordered separately.

- Order number GX26-3767 contains a binder and 3380 inserts.
- Order number GX26-3768 contains a binder and 3990 inserts.

Related Publications

These publications (not part of the Storage Subsystem Library) describe the IBM 3880 Storage Control models to which the 3380 can also attach:

- *IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*, GA26-1661
- *Introduction to IBM 3880 Storage Control Model 13*, GA32-0062
- *IBM 3880 Storage Control Model 13 Description*, GA32-0067
- *IBM 3880 Storage Control Model 23 Introduction*, GA32-0082
- *IBM 3880 Storage Control Model 23 Description*, GA32-0083.

Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage, GC26-4498, is intended for use with the 3380 subset of the SSL.

Other publications, referenced in this manual or providing additional related information, are included in "Bibliography" on page 145, which includes a short description of the relevant contents of each publication.

Summary of Changes

Second Edition, October 1988

Technical Updates

Technical additions have been made to the previously published material to describe:

- The MVS/DFP Storage Management Subsystem (SMS)
- The support provided for the 3380 in MVS/ESA operating environment.

Other Changes

This book has also been revised with minor technical updates. The book, *Cache Device Administration*, has also been added to the Storage Subsystem Library and an acronym list has been added to this book.

First Edition, September 1987

Technical Updates

Technical additions have been made to the original material to describe the following hardware and hardware features:

- IBM 3880 Direct Access Storage Models AJ4, BJ4, AK4, BK4, and CJ2
- The 3380 AJ4 and AK4 features #9341, #9342 and #9343
- The IBM 3880 Storage Control Models 3 and 23 features #3005 and #3010 that enable attachment of the 3380 AJ4 and AK4
- IBM 3990 Storage Control Models 1, 2, and 3

Library Structure

Some of the material in this manual was formerly contained in the following manuals, which are no longer available:

- *IBM 3380 Direct Access Storage: Planning and Use*, GC26-4208
- *IBM 3380 Direct Access Storage: Migration*, GC26-4197.

The reformatting of this material is part of an overall restructuring of disk storage documentation, as in Figure 2 on page xx

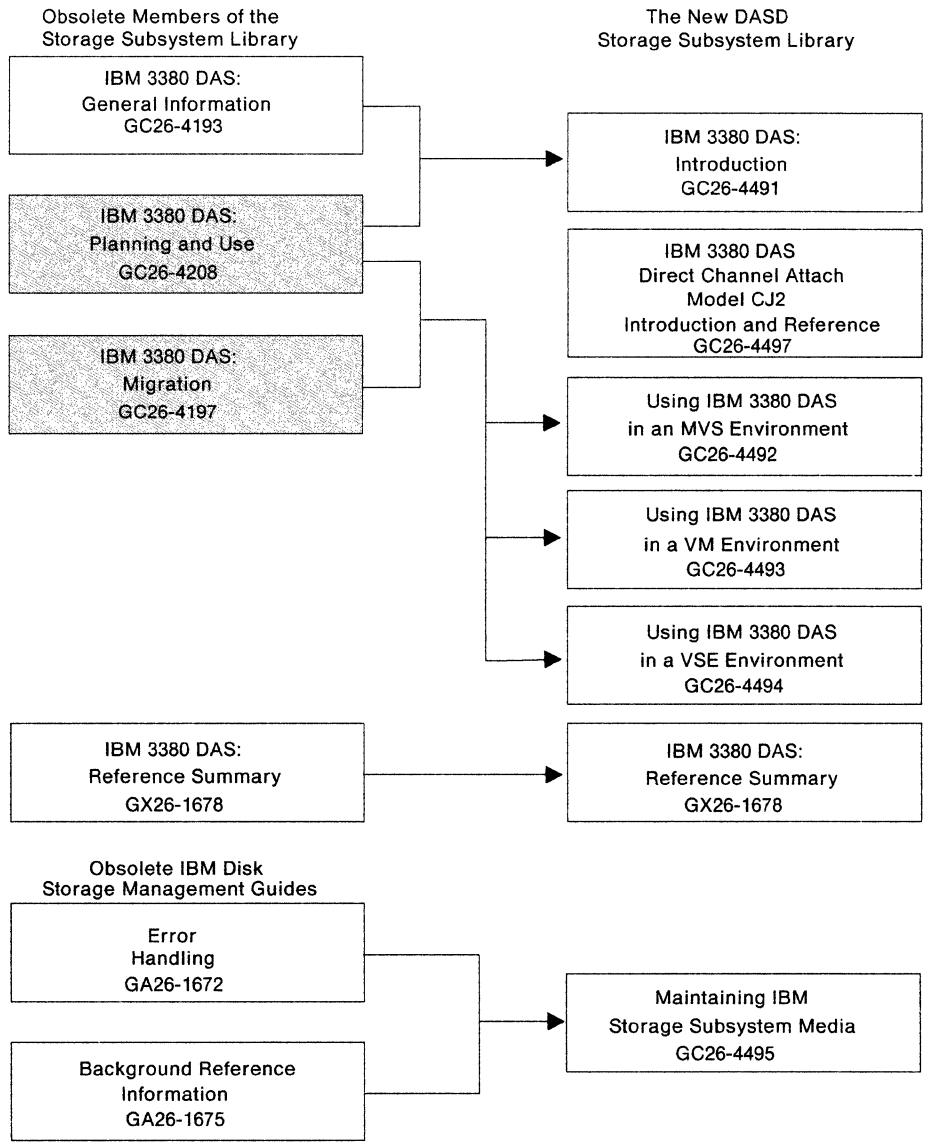


Figure 2. Storage Subsystem Library Restructure

Chapter 1. Introducing the IBM 3380 Direct Access Storage

The IBM 3380 Direct Access Storage is a high-speed disk storage device designed for online data storage. It offers quick and reliable access to data and is a cost-effective way to expand your online storage capabilities.

This chapter provides the following overview information for using 3380 disk storage in the MVS environment:

- A brief description of the 3380 family
- Hardware attachment requirements for 3380 models

3380 Direct Access Storage Overview

The twelve models of 3380 Direct Access Storage form four model groups. Several models offer increased storage capacity. However, all models of the 3380 have the same number of bytes per track and the same number of tracks per cylinder.

Volume Size	Standard Model Group	Extended Capability Model Group	Enhanced Subsystem Model Group	Direct Channel Attach Model
Single Capacity	A04 B04 AA4	AD4 BD4	AJ4 BJ4	CJ2
Double Capacity		AE4 BE4		
Triple Capacity			AK4 BK4	

Compared to the Standard models, the **Extended Capability** model group offers improved performance and availability and can provide concurrent transfer of data on two paths, from any two devices in the string. In addition, Models AE4 or BE4 provides twice the capacity of a Standard model.

The members of the **Enhanced Subsystem** model group are the most recent members of the 3380 family and provide further performance and availability enhancements. For example, the Enhanced Subsystem models have a faster seek time than other 3380 models, and when attached to the 3990 Storage Control Model 2 and Model 3, these 3380s can provide higher performance and availability for your data with concurrent data transfer on four paths. Furthermore, the triple capacity AK4 and BK4 models offer the 3380 family's lowest total cost of ownership on a per-megabyte basis.

The new 3380 **Direct Channel Attach Model** CJ2 provides both 3380 direct access storage and storage control functions in a single unit called a "C-unit"¹. The direct access storage functions of the Model CJ2 provide two paths for concurrent data transfer, and have improved seek characteristics over 3380 Standard and Extended

¹ The 3380 Model CJ2 contains two single-capacity volumes, and thus has half the capacity of other single capacity models, which contain four single capacity volumes.

Capability Models. The inclusion of a two path storage control function means that this 3380 model can be directly attached to a host processor channel.

As many as three 3380 B-units can be attached to an A-unit to form a string. As many as three Enhanced Subsystem B-units may be attached to a Model CJ2. The Model A04 provides a single path string, while the Model AA4 and Extended Capability models provide 2-path strings. The Enhanced Subsystem models can be configured as either 2-path or 4-path strings. A 4-path string consists of a minimum of two A-units, physically and logically connected, and can have as many as 3 B-units attached to each A-unit. The Model CJ2 and its attached B-units form 2-path strings.

This manual describes how to use A-units, B-units, and C-units in an MVS environment. A complete description of the characteristics, features, capabilities, and string configurations of the standard, extended capability and enhanced subsystem model groups can be found in *IBM 3380 Direct Access Storage Introduction*. For similar information on the direct channel attach Model CJ2, see *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*.

Attaching the 3380 to Storage Controls and Processors

Figure 3 on page 3 shows how the 3380 can be attached to IBM processors through various models of the 3880 and 3990 storage controls. Note that the table reflects only those devices (3380 models, storage controls, and processors) currently supported in one or more of the MVS operating environments. For information on all the possible DASD-storage control-processor combinations, see *IBM 3380 Direct Access Storage Introduction*. For a list of 3380 models supported in various MVS operating environments, see "Planning for Programming Support" on page 7.

In some cases, it may be necessary to install features on your 3380s or storage controls. These features are explained in *IBM 3380 Direct Access Storage Introduction*.

String Type	IBM Storage Control Model Attachment	IBM Processors
A04	3880 Model 2 (Storage Director 2) 3880 Model 3	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 4361 Processor ³ System/370 Models 158, 158-3 System/370 Models 168, 168-3 9375 and 9377 Processors
AA4	3880 Model 2 (Storage Director 2) 3880 Model 3 (AA4 string(s) only)	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 4361 Processor ³ System/370 Models 158, 158-3 System/370 Models 168, 168-3 9375 and 9377 Processors
	3880 Model 3 (with AD4 or AE4 string)	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 4361 Processor ³ 9375 and 9377 Processors
	3880 Model 13, 23	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
	3990 Models 1, 2, 3 ⁴	4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
AD4 AE4	3880 Model 3	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 4361 Processor ³ 9375 and 9377 Processors
	3880 Model 23	Standard Processors ¹ 4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
	3990 Models 1, 2, 3	4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
AJ4 AK4	3880 Models 3, 23	Standard Processor ¹ 4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
	3990 Models 1, 2, 3	4381, 308x, and 3090 Processors ² 9375 and 9377 Processors
CJ2	Not required, attaches directly to channel	4381, 308x, and 3090 Processors ² 9375 and 9377 Processors

Figure 3. IBM Processors and Storage Controls Compatible with 3380 Strings

Notes to Figure 3:

- ¹ The standard processors include 3031, 3032, 3033, 4341, and the 3042 Attached Processor Model 2
- ² The 3090 Models 300E, 500E, and 600E run VM/XA or MVS/XA with the Processor Resource/Systems Manager feature. Non-XA operating systems can be run on these 3090 processors.
- ³ Supported on 3-megabyte channel
- ⁴ Attachment to a 3990 requires the Model AA4 have a serial number of 15000 or greater for 60 Hz units or X0300 or greater for 50 Hz units.

Chapter 2. Planning for Installation

Adding new devices to your system requires adequate planning to ensure the process runs smoothly. Device installation *can* be streamlined, but you must be willing to spend the time up front to develop a detailed project plan.

Many of the problems found during device installation are typically things that can be identified long before the device arrived on the loading dock. These kinds of “surprises”—inadequate cable lengths, missing prerequisite features, user program device dependencies, or insufficient software support—cause delays that inhibit both user and system productivity. With careful and comprehensive planning, you can reduce the surprises at installation time to a minimum.

The major stages of installation planning are described in this chapter:

- Physical planning
- Software planning
- Administrative planning

All of these stages should be addressed in your project plan.

Appendix A, “Sample DASD Installation and Migration Project Plan” on page 121 contains an installation planning section that you can use to track these activities.

Physical Planning

Before you install new 3380 devices, there are a number of physical planning activities that you should complete. Use the information about physical characteristics of the 3380 in either *IBM 3380 Direct Access Storage Introduction* or *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference* to help you identify requirements for:

- Floorspace, electrical power, cables, and air conditioning
- I/O addresses
- Prerequisite features

Identifying 3380 Physical Requirements

The 3380 physical requirements: floorspace, electrical power, cables, and air conditioning are documented in *IBM Input/Output Equipment: Installation—Physical Planning for System/360, System/370, and 4300 Processors* and *9370 Information System Installation Manual—Physical Planning*. A summary of 3380 physical characteristics and features can be found in Chapter 4, “Planning the Hardware Configuration” on page 29. Work with your installation planning representative to determine the physical requirements of 3380 installation.

Reserving I/O Addresses

You must assign I/O addresses to all 3380 devices you are going to install on your system. These I/O addresses must be defined to the 3380s, 3880s, and 3990s and to the operating system (and Input/Output Control Program when present) during the installation process. At this time, decide which addresses will be assigned to the 3380 units you will install. For details about how to specify I/O addresses for the 3380, 3880, and 3990 units, see *IBM 3380 Direct Access Storage Introduction*.

For information on specifying addresses for the 3380 Model CJ2, see *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*. For information about defining I/O addresses to an operating system, see "Defining the 3380 to the System" on page 65.

Adding 3380 Strings to the Operating System

Processors that implement System/370 Extended Architecture will require an IOCP generation prior to installing one or more 3380 strings. To add 3380 units or strings to a system that already includes a 3380 string, MVS operating systems require either an I/O device generation (iogen) or execution of MVS configuration program (MVSCP). See "Defining New Devices to MVS and the Channel Subsystem" on page 66 for more details.

Devices can be defined to the operating system (during system generation) before the devices have actually been installed. During IPL, uninstalled devices are automatically placed offline to the system.

Perform the IOCP generation and either iogen or MVSCP prior to the arrival of the 3380 units. This system generation is the appropriate time for establishing MVS storage pools.

You will find examples of adding 3380 devices for each MVS operating system in Chapter 6, "Installing the 3380" on page 65. For more details about the system generation process, see the appropriate MVS system generation manual listed in the "Bibliography" on page 145. manual.

Nondisruptive DASD Installation

In subsystems configured for 4-path strings, the 3380 Enhanced Subsystem models can be installed without disrupting the availability of other storage resources. Models BJ4 and BK4 can be added to an existing string, or a second Model AJ4 or AK4 4-path string can be added without disrupting the availability of the existing string. Since this function is handled totally within the DASD subsystem, any operating system, Extended Architecture or System/370, that is connected to a 4-path subsystem can take advantage of this capability.

To take advantage of the nondisruptive DASD installation capability, you must pre-assign addresses for the devices planned for future installation and predefine these addresses to the operating system. Defining future DASD additions in present system generations (sysgen's and IOCP gens) will avoid future changes and IPLs.

For more information on nondisruptive DASD installation, see *IBM 3380 Direct Access Storage Introduction* and *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

Identifying Prerequisite Features

Check with the service representative to find out whether any prerequisite features will need to be applied to the DASD or the storage controls in your storage subsystem. If there are features to be applied, you need to schedule time to do so.

For a description of prerequisite features and to determine to which devices these features must be added, see *IBM 3380 Direct Access Storage Introduction*, or the appropriate storage control manual.

Software Planning

Software planning requires an understanding of the required release levels for minimum support in the environment in which the devices will be installed. It includes ensuring the availability of required:

- Licensed programs
- Software tools
- Publications.

Planning for Programming Support

Before you can use your new devices, you need the appropriate system software to support them. Determine which products are required for your environment. "Programming Support Levels," below lists the different MVS environments and directs you to figures showing the minimum version and release levels of the required and optional IBM licensed programs supporting the various models of the 3380.

If you do not have the appropriate system software, order the required products so that they will arrive before the 3380 devices. Install and test new software products before delivery of the 3380s.

Programming Support Levels

The following list refers you to an appropriate figure according to the system program being used and by the 3380 model to which your 3380s attach:

MVS/XA: See Figure 4 on page 8.

MVS/ESA: See Figure 5 on page 9.

MVS/370: See Figure 6 on page 10.

OS/VS1: See "Support in an OS/VS1 Environment" on page 11.

Although these tables list only the minimum version and release levels required for support of the 3380 models, you should always use the most current level available to take advantage of the latest product enhancements.

Program temporary fix (PTF) order numbers are not shown in these figures. Consult your IBM systems engineer for any PTFs required to provide a base level requirement for support. For more information about PTFs, see "Plan to Apply Software Fixes" on page 15.

MVS/XA Programming Support

The 3380 is supported under Multiple Virtual Storage/System Product (**MVS/SP**) Version 2 which implements System/370 Extended Architecture. MVS/SP provides system program support while MVS/XA Data Facility Product (DFP) provides device support. This combination, called the MVS/XA environment, supports all models and functions of the 3380, including dynamic path selection.

The following table lists the minimum version and release levels of the licensed programs required for support of the 3380 in the MVS/XA environment. Consult your IBM marketing representative for applicable PTFs or features.

3380 Model	Attached to	MVS/SP Version	MVS/XA DFP Version	ICKDSF Version	EREP Version
A04	3880 Model 2 or 3	2.1.0	1.1.0 or 2.1.0	7.0	3.1.0
AA4	3880 Model 2, 3, or 13	2.1.0	1.1.0 or 2.1.0	7.0	3.1.0
	3880 Model 23	2.1.2	1.1.2 or 2.1.0	8.0	3.1.0
	3990 Model 1, 2, or 3	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2
AD4	3880 Model 3	2.1.1	1.1.2 or 2.1.0	7.0	3.1.0 with Feature 2
	3880 Model 23	2.1.1	1.1.2 or 2.1.0	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2
AE4	3880 Model 3	2.1.2	1.1.2 or 2.1.0	7.0	3.1.0 with Feature 2
	3880 Model 23	2.1.1	1.1.2 or 2.1.0	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2
AJ4, AK4	3880 Model 3 or 23	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2
	3990 Model 1, 2, or 3	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2
CJ2	Attaches directly to channel	2.1.2	1.1.3 or 2.2.3	9.0	3.3.2

Figure 4. Minimum Programming Support in the MVS/XA Environment

Ordering MVS/XA Licensed Programs: The licensed programs in Figure 4 may be ordered from your IBM marketing representative using the program numbers below.

Licensed Program	Program Number
MVS/SP JES 2 Version 2	5740-XC6
MVS/SP JES 3 Version 2	5665-291
MVS/XA DFP Version 1	5665-284
MVS/XA DFP Version 2	5665-XA2
ICKDSF	5655-257
EREP Version 3	5658-260

MVS/ESA Programming Support

The 3380 is supported under MVS/SP Version 3 which implements ESA/370 architecture. MVS/SP provides system program support while MVS/DFP provides device support. This combination is called the MVS/ESA environment and supports all models functions of the 3380.

The following table lists the minimum version and release levels of the licensed programs required for support of the 3380 in the MVS/ESA environment. Consult your IBM marketing representative for applicable PTFs or features.

3380 Model	Attached to	MVS/SP Version	MVS/DFP Version	ICKDSF Version	EREP Version
A04	3880 Model 2 or 3	3.1.0	3.1.0	7.0	3.1.0
AA4	3880 Model 2, 3, or 13	3.1.0	3.1.0	7.0	3.1.0
	3880 Model 23	3.1.0	3.1.0	8.0	3.1.0
	3990 Model 1, 2, or 3	3.1.0	3.1.0	9.0	3.3.2
AD4	3880 Model 3	3.1.0	3.1.0	7.0	3.1.0 with Feature 2
	3880 Model 23	3.1.0	3.1.0	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	3.1.0	3.1.0	9.0	3.3.2
AE4	3880 Model 3	3.1.0	3.1.0	7.0	3.1.0 with Feature 2
	3880 Model 23	3.1.0	3.1.0	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	3.1.0	3.1.0	9.0	3.3.2
AJ4, AK4	3880 Model 3 or 23	3.1.0	3.1.0	9.0	3.3.2
	3990 Model 1, 2, or 3	3.1.0	3.1.0	9.0	3.3.2
CJ2	Attaches directly to channel	3.1.0	3.1.0	9.0	3.3.2

Figure 5. Minimum Programming Support in the MVS/ESA Environment

Ordering MVS/ESA Licensed Programs: The licensed programs in Figure 5 may be ordered from your IBM marketing representative using the program numbers below.

Licensed Program	Program Number
MVS/SP JES 2 Version 3	5685-001
MVS/SP JES 3 Version 3	5685-002
MVS/DFP Version 3.1.0	5665-XA3
ICKDSF	5655-257
EREP Version 3	5658-260

MVS/370 Programming Support

The 3380 is also supported under MVS/SP Version 1, which implements System/370 architecture. MVS/SP provides system program support. MVS/370 Data Facility Product (DFP) provides device support. This combination, called the MVS/370 environment, supports all models and functional capabilities of the 3380, except dynamic path reconnection.

The following table lists the minimum version and release levels of the licensed programs required for support of the 3380 in the MVS/370 environment. Consult your IBM marketing representative for applicable PTFs or features.

3380 Model	Attached to	MVS/SP Version	MVS/370 DFP Version	ICKDSF Version	EREP Version
A04	3880 Model 2 or 3	1.3.0	1.1.0	7.0	3.1.0
AA4	3880 Model 2, 3, or 13	1.3.0	1.1.0	7.0	3.1.0
	3880 Model 23	1.3.3	1.1.1	8.0	3.1.0
	3990 Model 1, 2, or 3	1.3.5	1.1.2	9.0	3.3.2
AD4	3880 Model 3	1.3.0	1.1.0	7.0	3.1.0 with Feature 2
	3880 Model 23	1.3.0	1.1.1	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	1.3.5	1.1.2	9.0	3.3.2
AE4	3880 Model 3	1.3.0	1.1.1	7.0	3.1.0 with Feature 2
	3880 Model 23	1.3.0	1.1.1	8.0	3.1.0 with Feature 2
	3990 Model 1, 2, or 3	1.3.5	1.1.2	9.0	3.3.2
AJ4, AK4	3880 Model 3 or 23	1.3.5	1.1.2	9.0	3.3.2
	3990 Model 1, 2, or 3	1.3.5	1.1.2	9.0	3.3.2
CJ2	Attaches directly to channel	1.3.5	1.1.2	9.0	3.3.2

Figure 6. Minimum Programming Support in the MVS/370 Environment

Ordering MVS/370 Licensed Programs: The licensed programs in Figure 6 may be ordered from your IBM marketing representative using the program numbers below.

Licensed Program	Program Number
MVS/SP JES 2 Version 1	5740-XY5
MVS/SP JES 3 Version 1	5740-XYN
MVS/370 DFP Version 1	5665-295
ICKDSF	5752-VS2
EREP Version 3	5658-260

Support in an OS/VS1 Environment

The OS/VS1 operating system Release 7 and Basic Performance Extensions (BPE) Release 4 with DFDS Release 1.4 supports 3380 Model A04 and AA4 strings attached to 3880 Storage Control Models 2 and 3. If necessary, the Speed Matching Buffer for 3380 feature of the 3880 can be installed.

OS/VS1 does not support dynamic path selection functions. When a 3380 Model AA4 string is attached to two 3880 storage directors, two devices can be accessed at the same time.

Planning for Software Tool Availability

You use software tools to identify what data sets exist on your system, to determine their characteristics, and to reorganize them on your storage configuration. Installation planning for 3380 DASD includes making sure you have software tools that support your DASD. This section lists the minimum release levels of IBM licensed programs that support the various models of the 3380. If you find a software tool is not available at your installation, order it so that you can become familiar with it before your new hardware arrives.

Software tools available from IBM for use in the MVS environment and the 3380 tasks they will help you with include:

- Data Facility Data Set Services (**DFDSS**), helps you move data and perform backup and recovery
- Data Facility Hierarchical Storage Manager (**DFHSM**), provides space management, data backup and recovery.
- Interactive storage management facility (**ISMF**), a tool provided with DFP, helps you analyze and manage both data and volumes interactively
- Resource Management Facility (**RMF**), measures processor, channel, storage control, and device activity
- Cache RMF Reporter, collects statistics and provides reports on storage controls with cache
- System Management Facilities (**SMF**), measures the activity of individual data sets
- Service Level Reporter (**SLR**), produces reports from the information obtained by SMF or RMF

Figure 7 lists the minimum release levels of optional IBM licensed programs that support the 3380 in the MVS environment. These tools are described below. Information about their usage may be found in Chapter 3, "Understanding Your Current Configuration" on page 19. For complete reference information for these tools, see the appropriate licensed program documents, listed in "Bibliography" on page 145.

3380 Model	Licensed Program	Minimum Support	Program Number
A04 AA4 AD4	HSM Version 1 or DFHSM Version 2	Release 3.2	5740-XRB
	DFDSS Version 1 or DFDSS Version 2	Release 1.0	5665-329
	DFDSS Version 1 or DFDSS Version 2	Release 2.1	5740-UT3
	DFDSS Version 1 or DFDSS Version 2	Release 1.0	5665-327
	DFSORT	Release 8.0 ¹	5740-SM1
	RMF Version 2	Release 4.1	5740-XY4
AE4	DFHSM Version 2	Release 1.0	5665-329
	DFDSS Version 1 or DFDSS Version 2	Release 2.1	5740-UT3
	DFDSS Version 1 or DFDSS Version 2	Release 1.0	5665-327
	DFSORT	Release 8.0 ¹	5740-SM1
	RMF Version 2	Release 4.1	5740-XY4
AJ4 AK4	DFHSM Version 2	Release 2.0 ²	5665-329
	DFDSS Version 2	Release 2.0 ²	5665-327
	DFSORT	Release 8.0 ¹	5740-SM1
	RMF Version 2	Release 4.1	5740-XY4
CJ2	DFHSM Version 2	Release 2.0 ²	5665-329
	DFDSS Version 2	Release 2.0 ²	5665-327
	DFSORT	Release 9.0	5740-SM1
	RMF Version 2	Release 4.1	5740-XY4

Figure 7. Software Tool Support in the MVS Environment

Notes to Figure 7:

- ¹ Full support of the 3990 requires DFSORT Release 9 plus an SPE or a higher level.
- ² Consult your IBM marketing representative for the applicable PTF number. For information on SPEs or products not listed, see your IBM systems engineer.

Data Facility Hierarchical Storage Manager (DFHSM)

DFHSM improves storage utilization by automatically managing space and making data available using a hierarchy of storage devices. DFHSM manages space by keeping active data sets only on fast-access storage devices, freeing available space on user volumes by deleting eligible inactive data sets or by moving them to lower cost-per-byte devices (a DFHSM *migration* process). User allocation of migrated data sets are automatically retrieved (*recalled*) by DFHSM. DFHSM provides data availability by automatically backing up new and changed data sets and by dumping volumes, via invocation of DFDSS full volume dump.

For information on how you can use DFHSM to move data to new DASD, see "Alternate Methods to Put Data on 3380s" on page 62. For detailed information on

the capabilities of DFHSM, see the appropriate manuals listed in the “Bibliography” on page 145.

Data Facility Data Set Services (DFDSS)

DFDSS Version 2 is a licensed program you can use to move data between like and unlike devices², and to perform volume or data set backup and recovery operations.

In an MVS environment, DFDSS runs in two modes: system DFDSS and stand-alone DFDSS. The stand-alone version of DFDSS is limited to one function, restoring from a physical DFDSS dump tape.

The capabilities of the system version of DFDSS that are described are:

- Copy data sets between like and unlike devices.
- Copy full volumes between like devices.
- Dump of data sets or full volumes to a sequential data set; the target media can be tape, DASD, or a mass storage virtual volume.
- Restore a DFDSS-created sequential data set to like devices. You can restore data sets or full volumes.
- Restore a DFDSS-created sequential data set to an unlike device. This allows you to move data sets between unlike devices.

For volume copies between like devices, DFDSS copies data at high speed, moving an entire track in a single I/O operation. For data set copies between like devices, most data sets are also copied at track speed. When copying data sets between unlike devices, DFDSS may invoke standard MVS utilities (IEHMOVE, IEBCOPY, IEBSAM, or IDCAMS) to move the data; for this reason, it may not copy data sets whose organizations these utilities do not support.

The data sets to be copied by DFDSS can reside on one or more DASD volumes and can be selected, through what is known as “filtering.” DFDSS also contains a “no-run” option, which allows you to find out which data sets are selected by the filtering (selection) process without actually moving the data. If desired, DFDSS can be invoked from a batch job, called by a user exit, or invoked from the interactive storage management facility (ISMF) as described below. For more information on DFDSS, see the DFDSS manuals listed in the “Bibliography” on page 145.

Data Facility Sort (DFSORT)

DFSORT is a high-performance data arranger that performs sort, merge, and copy processing of data sets. DFSORT processes EBCDIC, ISCII/ASCII, and user-defined data types. While processing data sets, user-written routines can perform record-level editing operations such as deleting, inserting, and altering records.

² “Like” devices have the same track capacity, and the same number of tracks per cylinder; that is, like devices can differ only in the number of cylinders per volume. “Unlike” DASD devices have different track capacities, and/or a different number of tracks per cylinder. For example, all members of the 3380 family are like devices.

All levels of DFSORT provide support for all models of the 3380 and the 3990. However, to take full advantage of the capabilities of the 3990, an SPE is required for DFSORT. For more information, see your IBM systems engineer.

For a general description of DFSORT, see *DFSORT General Information*. For detailed information, see *DFSORT Application Programming Guide*.

Interactive Storage Management Facility (ISMF)

ISMF provides interactive support to help you perform storage management activities using MVS/XA DFP Version 2 Release 2 (or later), or MVS/DFP Version 3 Release 1 (or later). ISMF helps you invoke the data set functions of DFHSM or DFDSS, or the edit functions of Interactive System Productivity Facility (**ISPF/PDF**).

ISMF offers a common, consistent interface to DFDSS and DFHSM through a panel interface. From these panels, you can use the data set application of ISMF to build lists of data sets using partially or fully qualified names and invoke DFDSS or DFHSM commands that move or copy data.

With ISMF, you can do the following without creating and debugging JCL or being concerned about storage management product control statements:

- If MVS/XA DFP Version 2 Release 2 and DFDSS Version 2 Release 2 are installed, line operators and list commands allow you to:
 - Copy data sets
 - Dump or restore data sets.
- If MVS/XA DFP Version 2 Release 3 and DFDSS Version 2 Release 3 are installed, line operators and list commands allow you to:
 - Copy data sets or volumes
 - Dump or restore data sets or volumes.
- If MVS/XA DFP Version 2 Release 2 and DFHSM Version 2 Release 2.1 are installed, line operators and list commands allow you to:
 - Backup and recover data sets
 - Migrate and recall data sets
 - Delete unwanted backup or migrated versions of data sets.

For a complete description of all the capabilities and functions of ISMF, see the *MVS/ESA* or *MVS/XA ISMF User's Guide*.

Resource Measurement Facility (RMF)

RMF, an IBM licensed program, measures the utilization and performance of system resources, including the processor, I/O devices, and storage. Resources, for measurement purposes, can either be in use or available over the period of time during which the system is processing a job. RMF gathers and reports the active time of various system resources and their users, and provides an analysis of contention for these resources.

RMF reports that can help you analyze performance are described in "Using RMF to Review I/O Workload" on page 22.

Cache RMF Reporter

Cache RMF Reporter consists of a user exit to the Resource Management Facility (RMF) and a Post Processor Report Program. The reports, useful in system tuning, provide statistics on storage controls with cache. The major benefits of Cache RMF Reporter are:

- Automatic data collection
- Cache statistics synchronized with other RMF reports
- Automatic collection of all hit ratios and other cache statistics
- Report format designed for review at a terminal.

For more information on using this tool, see "Tuning the System for Performance" on page 111 and *Cache RMF Reporter Program Description/Operations Manual*.

System Management Facilities (SMF)

SMF records I/O activity against individual data sets. You can extract the following information from SMF records:

- Average data set size for both tape and DASD
- Number of multivolume data sets
- Percentage of all data sets that are permanent
- Percentage of all data sets that are temporary
- Frequency of data set usage
- Average block size, block count, and EXCP count of each tape data set,

The analysis and reporting of this data can be done by either writing a user program or using the IBM Service Level Reporter (SLR). More information can be found in the SMF manuals listed in "Bibliography" on page 145.

Service Level Reporter (SLR)

SLR, an IBM licensed program, provides you with information about the applications and jobs in your installation. The source of information for the reports is SMF or RMF data, which have to be gathered. The sample programs provided with the product produce reports on the number of I/O operations. Release 2 of SLR can be used to obtain I/O counts by data set.

SLR collects information in one data base and presents it in a readily understandable format. SLR can merge data from several sources (for example, SMF and RMF) into one report. With graphics or reports that can be produced with SLR, you can gain a better understanding of volume activity, data set activity, and other aspects of your storage environment to determine where the heavily used data sets are and which volumes can be merged. More details about this report generator may be found in the SLR manuals listed in the "Bibliography" on page 145.

Plan to Apply Software Fixes

Work with your IBM account team to determine if there are any known problems that could affect your installation, and if so, plan to apply the fixes when you install the software. Applying fixes to the software before you install the hardware can save you time and problem-determination effort.

To assist in the installation of the 3380 device, a "hardware install index" is available from IBM Level One software support. This index contains general information for all operating systems, and specific information for each operating system and includes a list of software products that have been modified to support

the 3380 units and a list of pointers to existing Preventive Service Planning (PSP) information.

PSP information describes previously encountered problems that have a high potential impact, and refers you to Program Temporary Fixes (PTFs) available to correct these problems. The Field Maintenance Identifiers (FMID's) and PUT levels of your software products enable IBM Level One Support to tailor the PSP information to the release level of the software product(s) of concern. Remember to install these PTFs prior to the hardware installation.

Planning for Publications Availability

When you begin to perform the planning and moving data tasks, several IBM publications will be useful for reference. The "Bibliography" on page 145 lists publications pertinent to operating systems, licensed programs, and other topics. Order IBM publications in time to help you with the tasks ahead.

The Storage Management Library will be of particular interest to you. This library has been written to provide guidance on efficient storage management and to help you move toward system-managed storage. Consult the following manuals for up-to-date, task-oriented information:

- *MVS Storage Management Library: Configuring Storage Subsystems* describes how to configure storage hardware to balance performance, space utilization, and availability. It discusses the considerations for choosing storage hardware configurations and shows how to detect various performance and availability problems.
- *MVS Storage Management Library: Managing Data Sets* provides information on communicating with user groups, managing active and inactive data sets, managing catalogs and control data sets, establishing and enforcing data set policies, and providing data set security.
- *MVS Storage Management Library: Managing Storage Pools* describes how to plan for, design, implement and maintain storage pools. It also outlines the tasks that users must perform to take advantage of storage pools.

Your IBM representative can order the publications that you need.

Administrative Planning

In addition to physical and software planning, there are a number of administrative planning tasks that are vital to successfully installing devices and moving data. These tasks include:

- Creating a project team and project plan
- Verifying device delivery schedules
- Scheduling system time for volume preparation and data movement
- Training personnel
- Documenting new procedures.

Creating a Project Team and Project Plan

To handle the installation and migration process, you will want to establish an installation and migration team from among your data processing personnel. Work with your management and IBM representatives to determine the appropriate representatives for the device installation and migration team at your installation. Consider including:

- MVS system programmers or administrators
- Hardware configuration specialists
- Operations specialists
- Local tools specialists
- IBM service and account representatives
- User representatives.

It's also important to keep accurate records of your planning activities so that you can anticipate problems in scheduling device installation and data movement. A written project plan can help you determine what needs to be done to install and use the new devices. Make sure the project plan is available for all members of the device installation and data movement team to use as they complete their assigned tasks. The project plan should:

- Identify all tasks and responsibilities, including dependencies on other groups or events
- Assign direct responsibility for each task
- List target dates for completion of each task
- Identify meaningful checkpoints to evaluate progress
- Include statements of commitment from all involved groups.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a sample device installation project plan that you may want to use as a model.

Verifying Device Delivery Schedules

Contact your IBM marketing representative to verify that your new devices have been ordered, and to confirm the delivery date. Use this date to plan a tentative installation schedule for the 3380s. If you are ordering one or more storage controls for your new 3380s, verify that delivery will occur prior to the scheduled installation time.

Scheduling System Time for Volume Preparation and Data Movement

In most cases, you will need to schedule system time to prepare volumes and to move or recreate data. Obviously, the scheduling constraints of each installation will be different, but it's important to make arrangements well in advance. Negotiate with your user groups to determine the best time to prepare volumes and to move key data sets.

Training Personnel

Although installation of a new device should be transparent to most users, you will need to train some of your data processing personnel to use the new device. Operators will need to know how to power up a 3380 string and how to vary the devices online and offline; system support personnel may need to learn how to use available software tools to move data onto 3380 volumes.

You should plan hands-on training sessions for these people as soon as the 3380 units are available. For more information on the tasks that operators perform, see

Chapter 7, "Operating the 3380 under MVS" on page 77. For more information on moving data, see Chapter 8, "Moving Data in the MVS Environment" on page 87.

Documenting New Procedures

In addition to training your operators and system support personnel in the use of the 3380, you may need to document new operating and storage management procedures. These new procedures may include guidelines for:

- Data set allocation
- Data movement
- Data security
- Backup and recovery
- Volume initialization
- Media maintenance procedures.

Depending on your needs, you may want to document these policies and procedures as a user's guide, online briefing, or other type of presentation.

You should also plan to document the changes you make in your hardware configuration and in volume or data set placement. Any notes you make on migration problems, data set location and ownership, and application dependencies may provide valuable input for future device migrations.

Chapter 3. Understanding Your Current Configuration

To make the best use of your new hardware, you must understand how each device fits into your overall system configuration and how its distinctive capabilities can be used to best advantage. The key to effective use of the 3380 is to match its physical characteristics to the logical requirements of your data. Like hardware, data should follow a planned *configuration*; the way it is grouped, stored, and retrieved should depend on its logical data needs just as the placement of hardware depends on physical device needs. The topics discussed in this chapter help you understand your configuration and make effective use of the 3380:

- Determining Data Set Inventory
- Identifying Performance-Critical Data Sets
- Measuring the Effectiveness of your Current Hardware Configuration
- Sorting Current Volumes by I/O Activity
- Removing Device-Dependent Code

Before you introduce the 3380 into your MVS storage subsystem, take a comprehensive inventory of the data and storage in your installation. You can identify requirements for specific levels of performance, availability, and space by grouping data into categories and discussing these categories with your users.

If you are running under MVS/DFP Version 3, you can take advantage of the next step in efficient storage management, **system-managed storage**. With this approach, the system automatically manages data availability, performance, and space reclamation. System-managed storage also aids in device migration because users need no awareness of the physical device types in a storage group. Based on user requirements, the system can match the logical needs of the data to the physical characteristics of the installed devices, and ensure that each device is used most effectively. See Chapter 5, "Planning the Data Configuration" on page 47 for a more detailed discussion of system-managed storage.

As you move towards a more automated environment, you'll want to identify and eliminate device-dependencies in user programs that might constrain efficient use of your devices. Measure the effectiveness of your current hardware configuration by evaluating access paths and I/O activity. The more thoroughly you understand your current environment, the easier it will be to assess the impact of installing the 3380 on your data, programs, and users.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a checklist you can use to track the activities described in this chapter.

Determining Data Set Inventory

To move your data sets in the most logical and efficient manner and to ensure that you don't lose data sets, you must know what data sets are on your system, where they are located, and what their characteristics are. You can find that information in:

- Volume Table Of Contents (VTOC)
- Catalogs.

The VTOC contains entries for all data sets on a DASD volume. The VTOC entries describe the space on the volume and identify the location, extents, organization, and other characteristics of the data sets on the specific volume.

The second source of information about a data set is the catalog. Entries in catalogs typically contain: volume, security, ownership, extents, and/or allocation quantity information. The types of catalogs in an MVS environment are:

- Integrated catalog facility catalogs
- VSAM catalogs
- OS CVOLs.

MVS Tools for Determining Data Set Existence, Placement, and Characteristics

You can obtain information about a data set's location and size with the following tools:

- Interactive storage management facility (ISMF)
- IEHLIST utility
- Access method services LISTCAT command
- System Management Facilities (SMF) records and Service Level Reporter (SLR)
- Interactive System Productivity Facility/Program Development Facility (ISPF/PDF).

ISMF is an application of ISPF/PDF that assists you by generating lists of storage management information from the catalogs or the VTOCs on your system. Once generated, you can tailor these lists using ISMF's filtering, sorting, or entry-hiding capabilities.

The DFP function IEHLIST is a system utility used to list entries in an OS CVOL, the directory of one or more partitioned data sets, or an indexed or nonindexed VTOC. More information on this utility may be found in *MVS/XA Data Administration: Utilities*.

Access method services is a service program that is used with VSAM to establish and maintain catalogs and data sets. The access method services LISTCAT command can be used to list catalog entries in integrated catalog facility and/or VSAM catalogs. For more details, see the appropriate *MVS Access Method Services Reference* for integrated catalog facility catalogs or VSAM catalogs listed in the "Bibliography" on page 145.

A brief description of the information gathered by SMF and formatted by SLR is presented in "System Management Facilities (SMF)" on page 15 and "Service Level Reporter (SLR)" on page 15. For more information, see the appropriate manual listed in the "Bibliography" on page 145.

The utility option of ISPF/PDF provides a variety of functions that include library, dataset and catalog maintenance, displaying or printing VTOC entries for a DASD volume, and moving and copying data. For more information, see *Interactive System Productivity Facility/Program Development Facility: Program Reference*.

Identifying Performance-Critical Data Sets

After you have determined what data exists on your system, and gathered information about the way your data is presently being used, the next step is to identify the most performance-critical data.

This section lists typical performance-critical data sets. Be sure to check the current activity level of these data sets. If their activity level is high, these data sets may require careful positioning in your new DASD configuration. Alternatively, you may wish to place some or all of these data sets behind a cache storage control.

I/O activity to specific addresses on a DASD device (and thus to specific data sets) can be measured with Generalized Trace Facility (**GTF**). For information on how to use GTF in your environment, see the appropriate *MVS Service Aids* or *Initialization and Tuning Guide* listed in the "Bibliography" on page 145.

MVS SML: Managing Data Sets contains more information on critical data set placement.

Frequently Accessed System Data Sets

To achieve your performance goals, these data sets may require special placement, depending on their usage at your computer complex:

- Resource Access Control Facility (RACF) inventory data sets
- System log data set
- Paging data sets
- Swap data sets
- Spool checkpoint data sets
- Catalogs
- SMF data sets
- Procedure libraries
- SYS1.LINKLIB
- SYS1.COMDLIB
- SYS1.VTAMLIB.

Frequently Accessed DFHSM Data Sets

These data sets have an important effect on DFHSM performance:

- Backup control data set
- Migrate control data set
- Journal data set
- Off-line control data set.

Frequently Accessed IMS/VS Data Sets

These data sets might require careful positioning to optimize IMS/VS performance:

- Application control block libraries
- Format libraries
- Program libraries
- Message queues
- Write-ahead data sets.

Frequently Accessed DB2 Data Sets

The DB2 directory (DSNDB01) and the DB2 catalog (DSNDB06) might require manual positioning to enhance performance. The location of the catalog and directory can be changed with a recovery procedure as explained in *IBM DB2: System Planning and Administration Guide*.

Frequently Accessed User Data Sets

User data sets that are frequently accessed may require careful positioning to optimize system performance. Examples to check include CLIST libraries and procedure libraries. Your system may also have other frequently accessed user data sets.

Measuring the Effectiveness of Your Current Hardware Configuration

Part of understanding your current MVS environment is identifying limitations caused by inadequate access paths and inefficient volume utilization. There are tools available to help you evaluate the performance of your storage subsystem and identify the sources of I/O bottlenecks. Study a period of time that reflects the variability of the workload at your installation, and use the results to answer the following questions:

- Which devices have consistently high or low busy periods?
- Which channels have consistently high or low busy periods?
- Where are the peaks and valleys of I/O activity?
- Where is there contention for resources?

While devices with a high busy rate or high contention are not necessarily bad, these devices are prime suspects and should be investigated to see if they are causing poor I/O response times. When you find the contributors to poor I/O response times, you can place new 3380 devices and storage controls in your configuration to distribute the load more evenly and improve performance.

Using RMF to Review I/O Workload

The RMF reports described here can help you measure I/O utilization before and after modifying the I/O subsystem. RMF reports generated in the MVS/ESA, MVS/XA, and MVS/370 environments are listed, followed by pertinent measurement information contained in those reports. Different versions of RMF are used with MVS/370 as opposed to with MVS/ESA and MVS/XA; each version provides a similar report for I/O loading information. More information about how to use these reports may be found in *MVS SML: Configuring Storage Subsystems*. For reference and operation information, see the "Bibliography" on page 145 for the appropriate RMF publications.

- Channel loading

In MVS/ESA and MVS/XA, use the Channel Path Activity report.
In MVS/370, use the Channel Activity report.

Both reports provide a measurement of channel utilization (percent busy).

- Control unit loading

In MVS/ESA and MVS/XA, use the I/O Queuing Activity report.

This report provides logical control unit (LCU) measurements. Control unit utilization can be calculated by summing the connect times of the DASD attached to a storage director.

In MVS/370, use the Channel Activity report and the Direct Access Device Activity report.

The MVS/370 version of RMF does not provide a separate control unit report, but includes control unit measurements in both the Channel Activity report and the Direct Access Device Activity report.

Cache RMF Reporter can also be used to gather information on control unit loading. For more information, see *Cache RMF Reporter Program Description/Operations Manual*.

- Device loading

In MVS/ESA, MVS/XA and MVS/370, use the Direct Access Device Activity report

This report provides two measurements that directly indicate device loading: "Device I/O Rate" and "% Utilization." In MVS/370, "%Utilization" is called "% Busy."

RMF Measurements

The following RMF measurements from the channel and device activity reports can be used to determine the I/O path load:

- Percent channel busy and CPU wait can indicate where delays are occurring that might be contributing to processor wait time.
- Percent channel busy gives channel use and can indicate relatively over-used or under-used channels.
- Queue length for devices might indicate high wait times for satisfying requests on specific devices. The RMF device activity report provides the average DASD response time to service a request to devices on potentially critical channels. In addition, examine start I/O rate and control unit utilization when looking at access path activity.

Sorting Current Volumes by I/O Activity

Before combining old volumes on new DASD, obtain information about the I/O activity on the source volumes. Obtain several samples of the start I/O rate and the device utilization for the peak periods of system activity and non-peak periods (for example, second or third shift).

Next, rank the volumes by their utilization (busy) rate for each of the periods sampled. Volumes with significant variations between sample periods may require

special consideration, but the remainder can be sorted into high or low activity volumes.

Combine low-activity volumes if their activity is uniform, or if their occurrences of peak activity are relatively few. Medium-activity volumes can also be combined if each volume is used at a different time (for example, different shift, day, week, or month) or if the occurrence of peak activity for each volume is at a different time.

Initially, consider moving high-activity source volumes on a one-to-one basis. Exercise caution when combining high activity volumes with low-activity volumes, since the low-activity data may have a peak activity period that would severely impact the high-activity performance.

Higher channel and device utilization rates can be sustained without loss of performance, as a result of hardware functions such as dynamic path reconnect, device level selection, device level selection enhanced, and cache.

Discrepancies between the start I/O and the device utilization ranking may indicate data that is accessed with long channel programs. These volumes may tend to monopolize the data paths, and this should be considered when planning the new configuration.

Most systems have a few volumes of data that account for a high percentage of the I/O activity during one day or one sampling period. It is quite common for some volumes to have a high start I/O rate for a short period of time. For example, a volume might contain only programs or data used at the end of each month or each quarterly period. These volumes are usually low-activity volumes, and contain data that must be updated quickly to meet processing deadlines for critical applications. Identifying these performance-critical data sets and their location will help you make wise decisions if you choose to move them to new volumes.

Variations in the ranking of I/O activity between sampling periods may provide an opportunity for balancing low activity and high activity. However, application execution requirements may vary from shift to shift and day to day and therefore require future readjustments in data placement.

MVS SML: Configuring Storage Subsystems. contains more information on this topic.

Removing Device-Dependent Code

The removal of device-dependent code from application programs, job control language, procedures, and CLISTS will facilitate device installation and promote more efficient device utilization. Identifying where device-dependent code exists early in your planning cycle will let you make better decisions (like when and where to move it or how to remove it).

This section shows you where to look for device-dependent code, methods for locating it, and techniques that will eliminate some of these obstacles to moving data.

Where to Look for Device-Dependent Code

The following areas may have device-dependent information:

JCL, CLISTs, and Dynamic Allocation Lists

JCL DD statements should be checked for device-dependent parameters. Notify users before moving their data sets so they can search for and remove device-dependent code from the following areas:

DCB parameters. Use block sizes that provide efficient utilization on all your DASD and remove hard-coded block sizes in user programs.

Catalog all data sets. Catalog all data sets. Use esoteric allocation for new data sets. Remove Stepcat and Jobcat DD statements from JCL.

Master Catalog UNIT/VOLSER for the system residence volume. To eliminate device dependency, specify VOLSER = ***** and unit = 0000 in the master catalog for the system residence volume.

Generic UNIT specification, such as "3350." These need to be changed as appropriate, preferably to an esoteric name (such as SYSDA).

Accounting Routines

Accounting routines may be sensitive to the number of bytes per track or the number of tracks per cylinder and may require modification. The basis for charging for the use of DASD may also need to be changed.

EXCP Programming

Programs written at the EXCP level may need modification to run on 3380 devices, and may require careful scrutiny to eliminate device dependencies. Device dependencies commonly found here include track size and cylinder size, which on the 3380 are different than older DASD. Remember that the 3380 does not support track overflow. Also, the DEVTYPE macro cannot be used to determine the number of records per track. The TRKCALC macro can be used to calculate the number of fixed-length records that will fit on a 3380 track.

User Exits (JES, SMF, DADSM, OPEN)

Verify that user exits (JES, SMF, DADSM, and OPEN) are not sensitive to device characteristics, such as bytes per track, tracks per cylinder, or device type.

ISAM Random Processing

If you still use ISAM, this is a good time to convert to VSAM. If you do not convert to VSAM, remember that one of the facilities available to assist in random processing of indexed sequential data sets is that of retaining the highest level index in storage. In most cases, this will be a cylinder index, because a master index is not generated unless specifically requested. Unless the 3380 data set block size used is larger than 2048, the number of cylinders used on 3380 devices will be greater than that on 3350 devices. This means the cylinder index will be larger and will therefore require more storage. To compensate for this increased storage, it may be necessary to increase the size specified in the REGION parameter.

Identifying Device Dependencies

This section shows you ways to identify device dependencies at your installation.

Monitor DCB information

DFP provides a user exit for OPEN, which gives access to the user's DCB and can be used to check for hard-coded parameters. (For specific information about coding this kind of exit routine, see the appropriate *MVS Data Administration Guide*.)

Alternatively, you can code a JES exit to gain access to hard-coded DCB information. Exit 6 in JES2 and CSECT IATUX03 in JES3 permit inspection of JCL after it has been resolved and interpreted. The JES exits do not give access to dynamically allocated or TSO-allocated data sets. If you use the JES exits, code your exit routines to preselect the information that is to be recorded; this will help avoid the accumulation of large amounts of unwanted data.

For information about enforcing data set policy with user exits, see *MVS SML: Managing Data Sets*.

Scan JCL

Scan all JCL to ensure that DCB parameters are valid and appropriate for new devices. This can be done with a user-written CLIST or program.

Inspect Application Programming

Each Model AK4 and BK4 device has 39,825 tracks. In applications that do their own calculations in the following areas, half-word arithmetic should be done logically to ensure that the high-order bit is not used as a sign bit or that the calculation is not converted to full-word arithmetic:

- Extent arithmetic. Programs that do extent-arithmetic might encounter a data set with an extent greater than 32K tracks.
- Logical track number (TTR). The logical track number, as counted from the beginning of the volume or the beginning of an extent, can now be greater than 32K.

Techniques to Minimize Device-Dependent Code

The following programming practices will facilitate moving data to all devices in your storage subsystems, including 3380 devices. Note that some of these practices describe ways to eliminate device dependent code using the Storage Management Subsystem (SMS) component of MVS/DFP Version 3. Using SMS will help you move towards system-managed storage and achieve optimum utilization of DASD space.

- Catalog all data sets. In an SMS environment, all data sets must be cataloged.
- Remove explicit VOLSER and UNIT information from JCL. In a non-SMS environment, use esoteric names instead. With MVS/DFP Version 3, you can make the UNIT parameter optional for non-SMS managed data sets by specifying DEFAULT UNIT parameter.
- Use Automatic Class Selection (ACS) routines in an SMS environment to determine target volumes for new data set allocation. VOLUME and UNIT information is not required. See Chapter 5, "Planning the Data Configuration" on page 47 for a description of ACS routines.
- Use device-independent block sizes that make efficient use of track capacity.
- Use average block allocation in JCL statements for allocation of DASD space. See Chapter 5, "Planning the Data Configuration" on page 47 for a more detailed discussion of average block allocation.
- Avoid BLKSIZE specification on DCBs in application programs.
- Use the TRKCALC macro to calculate the number of records per track.
- Scan application programs and remove device dependencies.
- Use only relative block addressing.
- Do not create unmovable data sets.
- Convert ISAM data sets to VSAM data sets.

For more information on removing device dependent code and SMS see the *MVS SML: Managing Data Sets*.

Chapter 4. Planning the Hardware Configuration

Depending on your needs, you can configure a given set of devices in a variety of ways. Some configurations provide faster paging rates; some improve I/O rates; and others increase data availability. Whatever your goals, you want to make informed decisions about your hardware configuration so that the results of your efforts are predictable and appropriate for your installation.

The factors that you need to consider in designing an effective hardware configuration are:

- Hardware capability
- Capacity
- Performance
- Availability
- Shared DASD
- Cache.

This chapter discusses various aspects of hardware configuration to help you achieve more effective configurations. Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a checklist you can use to track these activities described here.

For a description of valid attachment configurations of the 3380 and IBM storage controls and more detailed information on 3380 functions, see the *IBM 3380 Direct Access Storage Introduction* or the *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*, or the *IBM 3990 Storage Control Planning, Installation and Storage Administration Guide*. For additional information on configuration, consult *MVS SML: Configuring Storage Subsystems*.

Understanding Hardware Capabilities

An understanding of the following hardware capabilities will help you plan the physical configuration of your 3380 devices and their associated storage control(s).

- Internal Paths
- Dynamic Path Selection (DPS)
- Device Level Selection (DLS)
- Device Level Selection Enhanced (DLSE).

Since you may have to sacrifice one configuration goal to achieve another, knowledge of hardware functions is essential to make informed decisions and sensible trade-offs in designing a configuration for your installation.

Internal Paths

Each 3380 A-unit model (except A04) contains two controllers, and each controller has four paths for accessing the devices on the string. Each successive 3380 model group provides improved access path capabilities between the controllers and the devices on the string.

Standard Model Internal Paths

The 3380 Model AA4 has four internal data paths attached to the two controllers by means of a 2x4 switch array. Since Model AA4 controllers have dedicated paths to specific devices, it is important to spread high-activity volumes across the available paths. At any instant, only one controller can use an internal path to access any specific device on that path. Simultaneously, the other controller can transfer data to or from any of the other 12 possible devices using another internal path. You can maximize performance on AA4 strings by placing high-activity volumes on different paths.

An AA4 string consisting of only four devices (a single 3380 AA4 unit) has only two internal data paths; path 0 with devices 0 and 1, and path 1 with devices 2 and 3.

Figure 8 shows which devices on an AA4 string share internal data paths.

INTERNAL PATH	DEVICES															
	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	S	S							S	S						
1			S	S							S	S				
2					S	S							S	S		
3							S	S							S	S

Figure 8. 3380 AA4 Internal Data Path Sharing. In this table, an S below a device indicates that it is on the internal data path listed at the left side of the chart.

Extended Capability Model Internal Paths

Extended Capability strings provide internal pathing from the controllers to the devices on the string that allows any device on the string to be selected by either controller on any of four internal paths. Thus, when one device is busy, all other devices in the 2-path string remain accessible to the other controller via any of three remaining internal paths. *Any two* devices in the string may be selected concurrently, one by each controller. The 2-path capability that is provided with these models provides Device Level Selection (DLS) capability; that is, any two devices may read or write data simultaneously. See "Device Level Selection (DLS)" on page 32 for additional information.

Note that when 2-path strings of Extended Capability models are intermixed with a 4-path Enhanced Subsystem string on a 3990 Model 2 or 3 Storage Control, the storage subsystem runs in DLSE support mode, and the 2-path strings have DLS capability.

Enhanced Subsystem Model Internal Paths

Enhanced Subsystem strings provide internal path capabilities that allow any controller on the string to access any device on any of four internal controller paths. The concurrent device access capabilities for Enhanced Subsystem models depend upon how the units are configured.

- When Enhanced Subsystem models are configured as **2-path strings**, the pathing capabilities are similar to those of Extended Capability models. There are two controllers that can access any of the devices (as many as 16) on any of the internal paths. *Any two* devices in the string may be selected concurrently, one by each controller. The resulting capability is DLS; see “Device Level Selection (DLS)” on page 32.
- When Enhanced Subsystem models are configured as **4-path strings**, there are four controllers that can access any device on the string (as many as 32) on any of the internal paths. *Any four* devices on the string can be accessed simultaneously, one by each controller. This capability is called DLSE; see “Device Level Selection Enhanced (DLSE)” on page 32.

The appropriate configuration feature must be used with the Enhanced Subsystem A-units as explained in *IBM 3380 Direct Access Storage Introduction*.

Dynamic Path Selection (DPS)

Dynamic path selection (DPS) is based on the DASD string having two controllers providing data transfer paths from the 3380 string to the storage directors, with selection of an alternate controller if one controller is unavailable. The storage directors, in turn, attach to two or more channels, thereby achieving multiple paths for transferring commands and data.

For system-initiated communication, a second I/O operation to the string can always be started, provided the two devices are on separate internal paths within the string. If the 3380 controller designated in the I/O address is busy or inoperative, the operating system or channel subsystem making the request is notified and can then select a path to the other controller to access 3380 data. The DPS function allows an alternate path to be established through another storage director and the other controller.

The DPS functions are:

- Simultaneous data transfer. Any two devices can transfer data simultaneously to or from the controllers, provided the devices are on separate internal paths within the string.
- Alternate controller selection. Each controller in a 3380 A-unit is capable of accessing any device in the 3380 string. With alternate path selection in a 2-path string, if one controller becomes unavailable or busy, another path is selected through another storage director to the other controller. No manual intervention is required. In a 4-path string, there are four controllers that can access any device in the string, providing even greater availability. An alternate controller can handle I/O operations for any device in the string (although only one device may transfer data at a time on a path).
- Volume reserve (System-Related Reserve). With system-related device reserve (on 370/Extended Architecture hosts only), a 3380 device can be reserved for use by a group of paths rather than a single path. Only the paths identified as part of the named group can access the device. System-related device reserve can be used to achieve higher levels of availability and I/O performance. It can also prevent simultaneous attempts to update critical data (for example, VTOCs).

Because VM/SP and VM/SP HPO do not support dynamic path selection (DPS), they do not accept commands associated with path group identification (such as SET PATH GROUP ID). If used, these commands cause a unit check condition with command reject in the sense data. Guest operating systems under VM interpret this to mean that dynamic path selection is not available.

For more information, see *Using the IBM 3380 Direct Access Storage in a VM Environment*.

- **Dynamic path reconnect.** When data is not being transferred (such as during a seek operation), the paths to the 3380 device are disconnected. During this time the path is often used to access another device. When the original device requires the path again, it must be reconnected.

On 370/Extended Architecture and ESA/370 hosts only, dynamic path reconnect allows the device to reconnect to the first available path within the group, rather than waiting for the use of the original path. This improves I/O performance for 3380 devices by reducing channel path access time, thereby reducing queueing time.

Device Level Selection (DLS)

Device level selection (DLS) is a capability of 3380 Models AD4, AE4, AJ4, and AK4 strings that provides two data transfer paths, one from each controller, to each device in the 2-path string (as many as 16 addresses). DLS extends the capabilities of DPS by allowing **any two** devices in the 2-path string to read or write data simultaneously.

When attaching 2-path AJ4 or AK4 strings to a 3990 Storage Control, DLS support mode must be set by the service representative at hardware installation time. For more information on 3990 DLS support mode, see *IBM 3990 Storage Control Introduction*. The Enhanced Subsystem A-units must have the appropriate feature as explained in *IBM 3380 Direct Access Storage Introduction*.

DLS uses two storage directors, in either 3880 or 3990 Storage Controls, that provide two data paths to each device in a 2-path string. When attached to the 3990, each of the storage directors operates as a single-path storage director.

3380 models with DLS offer improved data availability and overall system performance when compared to Standard models. If the selected device is not busy, it may be accessed even if another device on the 2-path string is reading or writing data. When one device on the 2-path string is busy, any of the remaining devices can be selected. This can reduce the amount of time an operating system needs to wait for a path to a device to become available.

Device Level Selection Enhanced (DLSE)

Device level selection enhanced (DLSE) is a capability of 3380 Models AJ4 and AK4 (requiring two interconnected A-units) that provides four data transfer paths, one from each controller, to each device in a 4-path string (as many as 32 addresses). DLSE extends the capabilities of both DPS and DLS by allowing **any four** devices in a 4-path string to read or write data simultaneously.

A 4-path string attaches only to a 3990 Storage Control Model 2 or 3, with the storage directors operating as multipath storage directors to provide four data transfer paths. The 3990 operates in DLSE support mode, set by the service representative at hardware installation time. For more information on 3990 DLSE support mode, see *IBM 3990 Storage Control Introduction*. The Enhanced Subsystem A-units must have the appropriate feature, as explained in *IBM 3380 Direct Access Storage Introduction*.

3380 models with DLSE offer improved data availability and overall system performance when compared to both Standard and Extended Capability models. If the selected device is not busy, it may be accessed even if any three other devices on the 4-path string are reading or writing data. For example, when heavy batch sequential applications or dump/restore activity tie up a path from a channel to a device, DLSE reduces disruption in throughput performance to the remaining devices in the 4-path string. End-user response can be more consistent during heavy workload periods with DLSE.

With DLSE, any one of the four paths can be quiesced (set offline to the host) without disrupting availability of the devices on the other paths. This allows additional B-units to be added to the 4-path string without disrupting availability of the existing devices, or addition of another 4-path string to the 3990 storage control without disrupting use of an established string. See *IBM 3380 Direct Access Storage Introduction* for additional information on the nondisruptive DASD installation capability associated with DLSE.

Speed Matching Buffer for 3380

The Speed Matching Buffer for 3380 (**SMB**) is a feature that can be installed on the storage directors of a 3880 Storage Control Model 2 or 3. SMB allows you to attach 3380 Model A04 and AA4 strings to processor channels with data transfer rates less than 3.0 MB per second.

Since an AA4 string attaches to two storage directors (preferably in different 3880s), **both** storage directors to which the AA4 string is attached must have SMB installed. Furthermore, if one storage director in a 3880 Model 3 has SMB installed, the other storage director in that 3880 must have SMB installed.

Note: SMB is not available for the other models of the 3380. The feature can be removed from the 3880 by a service representative if attachment of other models is necessary. Once SMB is removed from a 3880, the 3380 devices attached to that 3880 will no longer attach to processor channels with data transfer rates less than 3MB/sec. Feature 3005 or 3010 must be added to the 3880 to support AJ4 or AK4 strings.

In developing your hardware configuration, plan to isolate 3380 A04/AA4 strings on their own storage controls if they require the SMB feature. For more information, see *IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual*.

Access Arm Movement

About 50% of all seek and seek cylinder channel commands set the 3380 access arm in motion and cause the device to disconnect from the access path until the access motion completes (the other 50% reference the same track or cylinder). Seek time represents the time needed to traverse the cylinders between the access arm's present location and the target cylinder.

Average seek time is obtained by moving the access arm from each individual cylinder to every other individual cylinder and then taking the average move time for all combinations. As such, it is not completely valid by itself for comparing devices that have different storage capacities. If one device seeks across 200 megabytes of data in the same time that a second device seeks across 400 megabytes of data, although it may have a longer average seek time, the second device has greater performance potential for equal data capacity.

Channel Switches

Channel switches in your hardware configuration will increase the availability of data. Two-channel switch, two-channel switch pair additional, and eight-channel switch allow the two storage directors of a 3880 to attach two, four, or eight channels, providing as many as eight channel paths from the same or different processors to each 3380 string attached to a 3880. When you attach a 3380 string to two storage directors in different 3880s, each with eight-channel switch, data in the string can be accessed by as many as 16 different channels.

Note: A 3380 string **cannot** be attached to two storage directors in a 3880 if both storage directors are attached to the same channel (unless one of the identical channels is always offline when the other is online). The two storage directors attached to a 3380 string must be attached to different channels.

Each 3990 Storage Control can attach to as many as four channels. Four-channel switch, additional, is an optional feature that allows connection to an additional four channels. When 3380 strings are attached to 3990 storage controls, with four-channel switch, additional, the data on these 3380s can be accessed by as many as 16 channels. The channels may be on the same or different processors.

The 3380 Model CJ2 attaches directly to as many as two channels. These channels may be on the same or different processors. For more information, see *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*.

For more information about channel switching and the 3380, see the appropriate storage control manual listed in the "Bibliography" on page 145.

Capacity Considerations

Although an elaborate capacity planning study is not required to use 3380s, you should have a high-level understanding of space utilization in your current hardware configuration before installing 3380s. By understanding your current space needs, you can plan for orderly growth and reduce those performance problems that accompany insufficient space.

Figure 9 summarizes the data capacity of the various 3380 models.

Data Capacity	A04, AA4, B04, AD4, BD4, AJ4, BJ4	AE4, BE4	AK4, BK4	CJ2
Per Device	630 MB	1 260 MB	1 890 MB	630 MB
Per Unit	2 520 MB	5 040 MB	7 560 MB	1 260 MB

Figure 9. Summary of Data Capacity for 3380 Models

If you are replacing 3350s with 3380s, you can use the following space capacity guidelines to determine the number of 3380 volumes you will need.

- Slightly less than two full 3350 volumes can be placed on one volume of a 3380 Model A04, B04, AA4, AD4, BD4, AJ4, BJ4, or CJ2.
- Slightly less than four full 3350 volumes can be placed on one volume of a 3380 Model AE4, or BE4.
- Slightly less than six full 3350 volumes can be placed on one volume of a 3380 Model AK4 or BK4.

Include spare volumes in your capacity plans. If a device failure occurs, recover the data onto the spare volume. For more information on capacity planning, see *MVS SML: Configuring Storage Subsystems*

Performance Considerations

To meet your response time and throughput requirements (performance requirements), you must place your new devices at the appropriate points in the system configuration. Identifying your performance-critical data sets, as described in Chapter 3, "Understanding Your Current Configuration," and determining the number of high-activity volumes required will help you design a configuration that will satisfy not just the performance requirements, but all the requirements of your computer complex.

Response and throughput of disk storage operations are determined by the I/O per second requirements of the data and by the I/O per second capability of the devices and paths. Some of the factors affecting response and throughput are:

- Configuration
- Storage control selected
- Cache
- Number of paths available to a device
- Number of devices in the string
- Models of devices on a string
- Placement of data sets on specific devices
- Amount of data that can be transferred during one disk rotation (track capacity is 47 476 bytes)
- Amount of data that can be accessed without access arm movement (cylinder capacity is 712 140 bytes)
- Time needed to transfer all the data on a track or on a cylinder plus an additional amount for seek time and rotational delay
- I/O queueing time.

I/O Activity

Configuration guidelines with respect to the I/O activity level of a given volume include:

- Alleviate path-contention problems in most environments by configuring your DASD in 4-path strings attached to 3990 storage controls.
- Use a storage control with cache, this may provide significant response and/or throughput gains.

When there is more performance-critical data than will fit on the designated volumes, consider the time windows within which the various data sets are used. Combine several types of performance-critical data on the same volume or access path if the times during which they are used do not overlap significantly. An example of this is combining batch data with online data (the online data is used during the prime shift and the batch data is used after the prime shift).

Balanced I/O

When planning the configuration, attempt to maintain a balance within the input/output subsystem. Balance, in this case, means that the components are appropriately sharing the load. No channel, storage director, storage control module, or device should be so busy that it causes excessive delays. In a balanced subsystem, there is no single component that is the limiting bottleneck.

Peak Loads

The activity to a given device or data set will show variations, depending on both the time interval used for the measurement and the time of day at which the measurement is taken. If the configuration is to provide a certain level of performance, then peak load periods must be taken into account. In critical situations, peak loads are one of the major factors that determine your system configuration.

Skew

For a given configuration of strings attached to a storage control, the total load is seldom spread uniformly over the available access paths and devices. It is not unusual to find that 40% to 60% of the devices carry 90% or more of the total load. This nonuniformity of I/O activity is called skew.

The devices of most concern are the devices toward which the actual activity is most heavily skewed. These devices are required to do more than their share, in terms of I/O activity, and, if they are unable to handle the load, it must be redistributed if performance objectives are to be met.

Device Reserve

In a shared DASD environment, device reserve on a 3380 device may cause greater serialization than on a 3350, because of the larger data capacity. For shared volumes that have shown a high reserve delay, consider using MVS Global Resource Serialization (**GRS**) or moving the data to reduce serialization.

GRS, a feature available in the MVS environment, operates with JES3 to provide the option to use global serialization at the data set level. Using channel-to-channel (**CTC**) connections, GRS reserves a data set for use only by the requesting processor. Concurrent access to other data sets on the same volume is allowed by any other processor sharing that volume. For more details, see *MVS Planning: Global Resource Serialization*.

Availability Considerations

Availability, for a storage subsystem, is defined as the degree to which a system resource is ready when needed to process data. Consider the availability requirements of your data. Protect data that is critical to your installation in case of a channel or storage control outage. Data of this type should be accessible through separate storage directors or storage control modules on different channels.

Sometimes availability problems in the storage subsystem are easy to spot: for example, when a channel path fails unexpectedly and alternate paths must be used. But you might also have pervasive availability problems without realizing their extent. Availability problems can contribute to poor response time.

Lack of availability can severely impact the users of your system. It is important to consider how your hardware configuration affects both system and data availability.

Some techniques you can use to configure for availability are summarized below.

- Establish multiple access paths to 3380 volumes. When an internal path is busy, data should be accessible via another internal path. DPS and DLS/DLSE offer significant availability improvements for paths between storage controls and devices. (See “Understanding Hardware Capabilities” on page 29 for more information on DPS and DLS/DLSE.)
- Establish alternate paths through separate storage directors in a 3990 storage control, or through separate storage directors in different 3880 storage controls (3880 Model 23s must be in a dual frame configuration). If a storage director fails, data will be accessible via the alternate path.
- Establish alternate paths to alternate processor channels, using channel switches on storage control. If a channel fails, data should be accessible via another channel. (See “Channel Switches” on page 34 for more information on switching.)
- Plan for an alternate processor through channel switches, and establish rights to DASD access during system generation. Define the devices to both processors during system generation, then vary them offline to the backup processor and online to the primary processor. If the primary processor fails, the alternate processor can be used as a backup until repairs are made.
- Keep spare volumes available to be used in case of a device failure. Many installations use their spare volumes to store temporary data until needed for hardware recovery purposes. Remember that your service representative may need two 3380 volumes to fix a problem on a single device, and you will need space to store the data until the problem is resolved. A common guideline is to keep one spare volume for every string.
- Consider using the 3990 Model 3 dual copy function to improve data availability. DASD fast write can be combined with dual copy, permitting accesses to see less of the overhead associated with dual copy. Dual copy is described in “3990 Model 3 Nonvolatile Storage, DASD Fast Write, Dual Copy” on page 40.

The manual *Component Failure Impact Analysis – An Availability Management Technique* describes how to plan and configure for hardware availability.

Shared DASD Considerations

Sharing DASD allows multiple systems to access data stored on common volumes. Although sharing improves job scheduling and data availability, you should be aware that it may result in device contention and performance degradation; several systems can be “locked out” while one system has the volume reserved for I/O activity.

Evaluate your use of shared DASD before planning changes to your hardware configuration. Consider the following general recommendations on sharing:

- Avoid sharing work volumes (for temporary data sets) among multiple systems. This will prevent jobs running on one system from scratching data sets on work volumes while those same data sets are in use on another connected system.
- Avoid sharing system volumes among multiple systems. System volumes contain many system data sets that are required to be unique to each system.
- If you are using volume pooling, and you want to share pooled DASD among systems, ensure that each system has the volumes mounted with different use attributes; for example, volumes mounted as public on one system should be mounted as private on other systems.
- Isolate shared volumes from dedicated volumes in the complex.
- If you share high-activity data sets, be sure to place them on the highest-performance devices in your installation and to plan for heavy I/O rates in your performance projections. These data sets (for example, RACF inventory and DFHSM control data sets) are good candidates to place under a cache storage control.
- JES3 can be used to handle some of the problems caused by sharing data. JES3 provides data set serialization on shared devices, maintaining data integrity by preventing simultaneous job scheduling of jobs that update the same data sets.
- GRS can be used to provide data set serialization on shared devices, to maintain data set integrity. GRS can also reduce device contention by changing volume reserves into data set enqueues.
- If you are using DFHSM in a shared DASD environment, you must determine which processor performs migration and backup functions for the shared DASD, and whether the catalogs for the data sets on these volumes are accessible to all sharing systems. Recovery procedures for shared data requires careful planning.

More information on shared DASD can be found in *MVS SML: Configuring Storage Subsystems*.

Guest System Considerations

Running MVS as a guest operating system under VM requires consideration of special configuration needs and capabilities. A thorough treatment of guest system hardware configuration is well beyond the scope of this book; see *VM Running Guest Operating Systems* for a more detailed discussion of this topic.

Consider the following general recommendations for hardware configuration in guest operating system environments:

- Dedicate devices and channels to guest systems. Because their heavy I/O activity will degrade performance for interactive users, MVS volumes should always be isolated from those managed by VM. For information on configuring strings of dedicated devices and channels on VM, see *Using the IBM 3380 Direct Access Storage in a VM Environment*.
- Remember that MVS guest systems use device reserve/release and string switching functions that are not otherwise available to VM. Configuration possibilities will vary accordingly.

Storage Control Selection

Three types of storage controls are available for 3380 attachment: direct channel attachment, non-cache, and cache. The 3380 Model CJ2 has a 2-path storage control function and attaches directly to processor channels. For detailed information on the 3380 Model CJ2, see *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*. Non-cache storage controls include 3880 Models 2 and 3, and 3990 Model 1 and Model 2. Cache storage controls include 3880 Models 13 and 23 and 3990 Model 3.

This section briefly describes how cache works, why you may need it, and what data sets you might want to put on the volumes attached to cache storage control. For an explanation of which 3380 models attach to which 3880 models and 3990 models, see *IBM 3380 Direct Access Storage Introduction*. For detailed information on storage controls, see the appropriate manual listed in the Bibliography.

Cache Storage Controls

The 3880 Model 13 and 23, and the 3990 Model 3 contain electronic storage called **cache**.

Storage Control	Available Cache Sizes
3380 Model 13	4 MB, 8 MB
3380 Model 23	8 MB, 16 MB, 32 MB, 48 MB, 64 MB
3990 Model 3	32 MB, 64 MB, 128 MB, 256 MB

Figure 10. Storage Control Cache Sizes

Cache acts as an intelligent, high-speed buffer to DASD, retaining frequently-read data for faster access by the processor. Having data available in the cache can greatly improve I/O response time. Access time between the cache and the channel is much shorter than that between the DASD and the channel; cache accesses eliminate the seek and rotational time inherent in DASD accesses.

For configuration purposes, 3880 cache storage controls can be attached to as many as two strings of 3380 devices, with as many as 16 volumes on each string. 3990 Model 3 storage controls can be attached to two 3380 AJ4 or AK4 4-path strings, or four 2-path strings, supporting as many as 64 volumes.

You can read about the types of data suitable for cache storage in *MVS SML: Configuring Storage Subsystems*. For more information on the use of cache storage, see the appropriate *Introduction* manual for your storage control:

3880	Model 13	<i>Introduction to IBM 3880 Storage Control Model 13</i>
3880	Model 23	<i>IBM 3880 Storage Control Model 23 Introduction</i>
3990	Model 3	<i>IBM 3990 Storage Control Introduction</i>

3990 Model 3 Nonvolatile Storage, DASD Fast Write, Dual Copy

The 3990 Model 3, which contains nonvolatile storage, (an area of electronic storage with battery backup), offers two significant new functions that can improve data availability and system performance:

- **DASD fast write** allows you to store data simultaneously in cache storage and nonvolatile storage. This is a "fast" write because once the data is written to nonvolatile storage, the system can continue processing without having to wait for the data to be written to DASD. At a later time the data is destaged from nonvolatile storage to DASD. This extends cache performance improvements to write operations. DASD fast write can be activated or deactivated on command.
- **Dual copy** allows you to maintain two identical sets of data on two different DASD volumes. The volumes are considered a duplexed pair, one primary (online) and one secondary (offline). As in a normal write operation, data is written simultaneously to cache and to the primary volume.

Because both DASD fast write and dual copy functions can be active at the same time, you can achieve significant improvements in performance and availability using these functions of the 3990 Model 3.

Nonvolatile storage is supported in Extended Architecture and ESA/370 environments and requires more than the minimum level of software support. For more information on nonvolatile storage, see *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

Configuring 3380 Strings

Your I/O performance and availability objectives (and the storage controls' attachment possibilities) will determine how many and what model 3380s to attach to your storage controls. Select the number of units and the storage control to meet your performance requirements. With the ability to handle higher I/O rates, and with cache, it is possible to mix data with different requirements on the same string and still achieve your performance goals for this data.

A set of volumes with a high I/O activity requirement and a limited capacity requirement might be configured as a 3380 Model AJ4 unit with a single BJ4 unit. A better solution to a high I/O activity requirement would be a storage subsystem with 3990 Storage Controls and a full 4-path string of Enhanced Subsystem 3380 models.

For more information on configuring storage subsystems, see *MVS SML: Configuring Storage Subsystems*.

Planning the Hardware Configuration

Now that you've reviewed the design considerations for your system service objectives, you're ready to explore various possibilities for building more effective hardware configurations.

There are many different ways to configure DASD, each with its own performance and availability characteristics. The examples presented in this section will show you a few ways to configure 3380s and storage controls to improve performance and/or availability. For a detailed discussion of possible configurations and power sequence cable diagrams, see *IBM 3380 Direct Access Storage Introduction*. By combining or adapting the configuration examples, you can plan a hardware configuration to meet the needs of your system.

The configurations shown in Figure 11 on page 42 through Figure 15 on page 46 assume that the 3380 A-units have two string controllers (Models AA4, AD4, AE4, AJ4, AK4 and CJ2).




Note: The connecting lines in these diagrams represent the CTLI cables; power sequence cables are not shown here.

For information on the features listed in the examples, see *IBM 3380 Direct Access Storage Introduction*.

Basic 3380-3880 Configuration

A basic 3380-3880 configuration is shown in Figure 11.

1-8 Channels



with Feature 9431

SD = Storage Director

A1-A2 = Controllers 1-2

Figure 11. Example of 3380 Model AE4 and AK4 Strings Attached to a 3880

In this example, two 3380 strings are sequentially connected to both storage directors of a 3880 Model 23. One string contains a 3380 Model AD4 controller followed by a mixture of BD4 and BE4 units; the other string contains a 3380 Model AK4 controller with a mixture of BJ4 and BK4 units. The 3880 storage control has the 3380 AK4 support feature installed. This configuration is also valid for a 3880 Model 3, which would require Feature 3005.

Intermixed Strings on 3990 Single-path Storage Directors


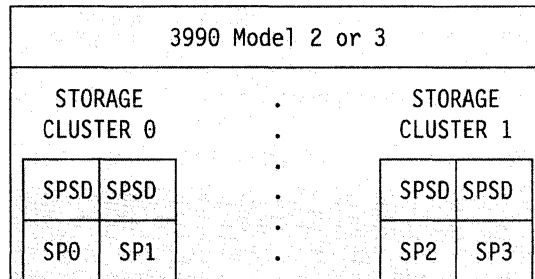
There are two rules for intermixing strings of different model groups on the same single-path storage director:

- One Model AA4³ string can be attached with one Model AD4 or AE4 string.
- One Model AA4, AD4, or AE4 string can be attached with one Model AJ4 or AK4 2-path string.

Figure 12 is an example of four 3380 2-path strings, with two sequentially connected to each of the paired storage paths in the same 3990 Model 2 or 3.

In this example, there are four 3380 strings, two sequentially connected to each of the paired storage paths in the same 3990 Model 2 or Model 3. Each string may contain a mix of models from the same model group.

As Many as 8 Channels As Many as 8 Channels



with Feature 9432 on AK4

SPSD = Single Path Storage Director .
 SP = Storage Path . = Power and service boundary
 A1-A2 = Controllers 1-2 .

Figure 12. Intermixed 3380 2-path Strings Attached to a 3990 Model 2 or 3

³ Attachment to a 3990 requires that the Model AA4 have a serial number of 15000 or greater for 60 Hz units or X0300 or greater for 50 Hz units.

Four-Path and Two-Path String Intermix

Attaching 2-path strings and a 4-path string to the same 3990 storage control is shown in Figure 13. The 3990 Model 2 or Model 3 is in DLSE mode, with 4 separate paths to the 4-path string and 2 separate paths to each 2-path string.

as many as 8 channels as many as 8 channels

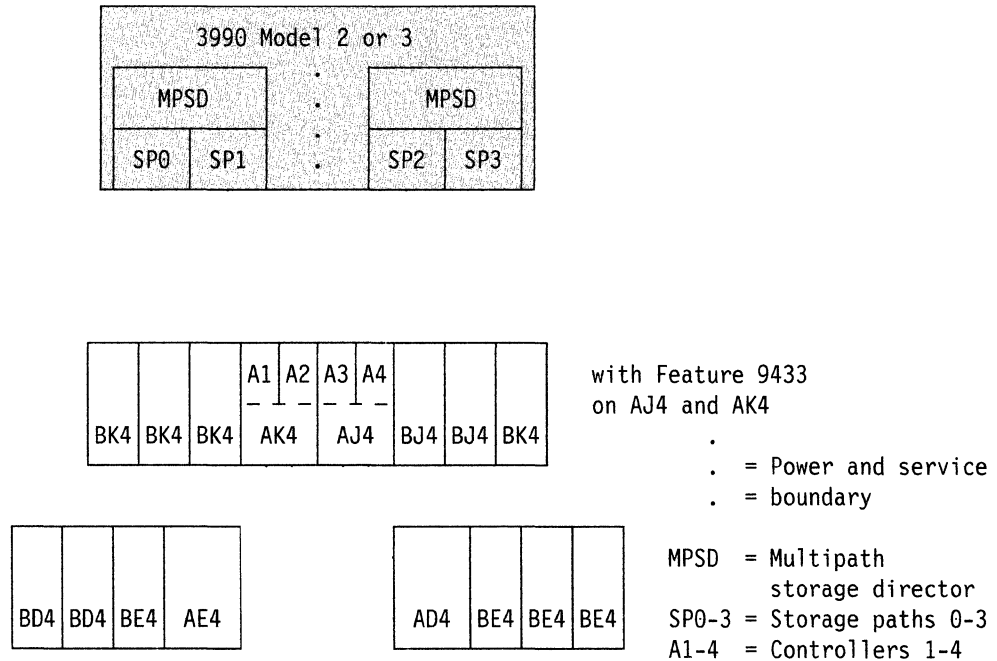




Figure 13. 3380 4-Path and 2-Path Strings Intermixed on a 3990 Model 2 or 3

The controllers (A1-A4) in the 3380 4-path string are sequentially connected to the storage paths (SP0-SP3) in the 3990 Model 2 or Model 3. The two controllers in each 2-path string attach to either storage paths 0 and 2 or storage paths 1 and 3, as shown in the example. Attaching 3380 strings to 3990 storage controls in a 4-path configuration requires the two A-units in the 4-path string to be logically and physically connected to each other.


A 3380-3990 4-Path Configuration for High Availability

Figure 14 shows two 3380 4-path strings sequentially connected to the paired storage paths in the same 3990 Model 2 or 3. This configuration provides high availability for the data on these volumes.

as many as 8 channels as many as 8 channels



. = Power and service boundary



MPSD = Multipath storage director
 SP0-3 = Storage paths 0-3
 A1-A4 = Controllers 1-4

with Feature 9433
 on all A-units



Figure 14. 3380 Model AJ4 and AK4 4-Path Strings Attached to a 3990 Model 2 or 3

As shown in this example, the two 3380 4-path strings must connect in sequence (A1-A4) to the storage paths (SP0-SP3) in the 3990 Model 2 or Model 3. Strings may contain mixtures of Enhanced Subsystem models. Attaching 3380 strings to a 3990 Model 2 or Model 3 storage control in 4-path configuration requires the two A-units in each string to be logically and physically connected to each other.

CJ2 Configuration

Attaching a 3380 Model CJ2 string directly to a processor channel is shown in Figure 15.

Channel Channel



SPSD = Single Path
Storage Director
SP = Storage Path
A1-A2 = Controllers 1-2

Figure 15. Example of 3380 Model CJ2 String Attached Directly to Channel

In this example, the 3380 Model CJ2 string contains a mixture of Enhanced Subsystem models. The Model CJ2 contains both a string controller and a storage control, and attaches directly to a processor channel.

Documenting the Hardware Configuration

After determining the configuration techniques, document the planned hardware configuration, including the rationale for choosing various configurations. A written hardware configuration plan should include:

- Diagrams of each new or changed DASD string, including storage control and processor attachments.
- A summary of capacity (number of devices/volumes) for each string.
- A description of the elements of the configuration that ensure hardware availability for each string: alternate paths, controllers, storage directors, storage controls, and channels.
- Diagrams showing location and attachments for shared DASD.
- Key contacts for hardware configuration.

Review the plan with all involved system support, operations, and user groups to ensure its viability.

Chapter 5. Planning the Data Configuration

If you are moving data sets from old DASD to new DASD, or adding DASD to your storage subsystems, you have to decide which data sets will be placed on the new volumes. This process includes identifying data sets to move, and data sets to create or allocate on the new volumes. For data sets that can be moved, this chapter identifies tools available to move them. Considerations on the placement of data sets are included here to help you plan your data configuration.

This chapter also discusses the use of storage pools because it will facilitate the task of storage management. With MVS/DFP 3.1.0, you can implement **system-managed storage**. System-managed storage describes the concept of letting the system take over storage management tasks previously handled by the storage administrator and users. This section provides more detail on how the system determines data set placement and automatically manages data availability and performance.

Once you evaluate the storage management requirements for your data processing installation and plan a configuration, you must document the plan. This chapter discusses why it is important to document the configuration and data movement plan and offers some guidelines on what information should be included in the plans. The importance of establishing special backup and recovery plans for use while moving the data is also emphasized.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a data configuration planning checklist you can use to track the completion of the activities described here.

Additional information on this topic is contained in the MVS Storage Management Library. See the appropriate manuals listed in the "Bibliography" on page 145.

Identifying Data to Move or Create

You can move most types of data sets with tools discussed in "Planning for Software Tool Availability" on page 11. Some data sets should be re-allocated rather than moved since they have allocated information (for example, SYS1.LOGREC). Re-allocating them will usually require changing JCL to define them on the new device.

Many of the important classes of data sets are described in the following sections, along with a discussion of considerations for moving or reallocating. Where it is possible to move data sets, tools that can be used to do so are listed.

Volume Table of Contents (VTOC)

VTOCs and indexed VTOCs are created on the new device using Device Support Facilities.

Considering the capacity of the 3380 devices, the use of indexed VTOCs becomes an important performance consideration. An indexed VTOC improves I/O performance by allowing direct access to the correct data set control blocks, eliminating lengthy sequential searches of the VTOC. The index also manages free

space information, reducing the number of I/O operations needed to obtain or release space on the volume.

For more information, see "Calculating the Size of VTOC and VTOC Index" on page 73.

Work Files

Work files used by sort programs, compilers, and applications are usually allocated using generic unit names. Generic unit names used for work files should be changed during system generation to associate them with the new devices.

Page and Swap Data Sets

Page and swap data sets are not moved to the new devices. These data sets are defined on the new devices with access method services and re-established during IPL. They are candidates for early assignment to 3380s.

JES2 Spool Data Sets

JES2 permits mixed device types as SPOOL volumes. JES2 supports the dynamic addition and deletion of SPOOL volumes; thus, by using operator commands (\$S SPOOL, \$P SPOOL, \$Z SPOOL) spool volumes may be added or deleted without interrupting JES2 processing.

Adding and Deleting Spool Volumes

A new spool volume may be started dynamically (via the \$S SPOOL operator command). The \$S SPOOL command may also be used to format the new volumes.

An existing spool volume may be drained and deleted (via the \$P SPOOL operator command). This will cause all work on the spool to be processed, and prevent any available space on the volume from being allocated. The volume will then become deallocated and can be removed.

Note: If ALL spool volumes are deleted, JES2 will not be able to perform work. Therefore, new volumes should be added before all existing volumes have been deleted.

Moving Spool Data to New Volumes

The Spool Offload facility can be used to offload spool data and delete a spool volume, and to reload the spool data to new spool devices. Any jobs and SYSOUT residing on the volume to be drained can be offloaded to an offload data set, and then reloaded. JES2 may also be cold started with new spool volumes after all jobs and SYSOUT have been offloaded. The offload data set may then be reloaded to the new spool volumes.

The BUFSIZE, TGSIZE, TGNUM, SPOOLNUM, and TRKCELL parameters on the SPOOLDEF JES2 initialization statement all have an effect on storage and device utilization. Generally speaking, the maximum BUFSIZE specification provides the best DASD and storage utilization.

See *MVS/Extended Architecture Operations: JES2 Commands* for detailed information about these commands. For information on the specification of parameters and their relation to maximizing device utilization and system performance, see *System Programming Library: JES2 Initialization and Tuning*

JES2 Checkpoint Data Sets — MVS/SP 2.1.5

In MVS/SP Version 2 Release 1.5 and prior releases, a JES2 Checkpoint data set can be allocated to a new device in two ways:

1. Use the Spool Offload facility to offload all jobs and SYSOUT in the JES2 complex. Then perform a JES2 cold start with the new device(s) used for the checkpoint data set(s), and reload the offload data set.
2. Use checkpoint duplexing to migrate as follows:

Perform a configuration-wide warm start of JES2 with the new device specified as the duplex (alternate) checkpoint volume (DUPLEX parameter on the CKPTDEF initialization statement).

Bring JES2 down cleanly (that is, stop all JES2 functions including devices, initiators, and jobs) prior to executing the next step. This ensures the alternate checkpoint data set has the most current data, and minimizes the possibility of disastrous errors as jobs are processed.

Perform a second configuration-wide warm start of JES2 with the new device specified as the primary checkpoint volume (PRIMARY parameter on the CKPTDEF initialization statement). Use the ALTCKPT JES2 initialization option to cause the duplex data set to be read at initialization, as the source of the checkpoint data. This data is then used to rewrite the primary checkpoint data set.

See *System Programming Library: JES2 Initialization and Tuning* for more information on the checkpoint data set, and data set placement in regard to JES2 performance.

Note: If you are doing a cold start using a new checkpoint data set, or doing a configuration-wide warm start using the ALTCKPT option, JES2 may issue the \$HASP479 message. In this case, the message is part of normal processing and does not indicate an error condition. This message is described in *MVS/Extended Architecture Message Library: JES2 Messages*.

JES2 Checkpoint Data Sets — MVS/SP 2.2.0

In MVS/SP Version 2 Release 2.0 and later releases, a JES2 Checkpoint data set can be allocated to a new device in two ways:

1. Use the Spool Offload facility to offload all jobs and SYSOUT in the JES2 complex. Then perform a JES2 cold start with the new device(s) used for the checkpoint data set(s), and reload the offload data set.
2. Use the Reconfiguration Dialogue, which allows the dynamic allocation of an alternate checkpoint data set on a new device without requiring a restart of JES2. This operator dialogue will guide the operator and/or systems programmer in dynamically modifying the checkpoint data set configuration (data set name and volume location).

The JES2 checkpoint data set configuration may be operated in either Primary/Duplex mode or in Dual mode. The checkpoint reconfiguration is permitted when operating in either mode, and allows for a replacement to be specified for each of the checkpoint data sets.

Use the \$T CKPTDEF, RECONFIG= YES JES2 command to enter into a dialogue in which the operator will supply the information needed to allocate an existing or create a new checkpoint data set. These data set(s) will be used as replacement(s) for the existing one(s).

You should be familiar with the reconfiguration dialogue before attempting to specify new checkpoint data set(s). For background information on the checkpoint configuration, see *System Programming Library: JES2 Initialization and Tuning*. For examples of the Reconfiguration Dialogue, see *MVS/Extended Architecture Message Library: JES2 Messages*. For more information on the \$T CKPTDEF command, see *MVS/Extended Architecture Operations: JES2 Commands*.

Note: When doing a configuration-wide warm start using the CKPT2 option in Duplex mode, JES2 may issue the \$HASP479 message. In this case, the message is part of normal processing, and does not indicate an error condition. This message is described in *MVS/Extended Architecture Message Library: JES2 Messages*.

JES3 Spool Data Sets

JES3 allows spool data sets to be allocated on mixed device types. The device types supported by any specific version and release of JES3 is documented in *MVS/Extended Architecture System Programming Library: JES3 Initialization and Tuning*.

The kind of JES3 start required to migrate a spool data set to a new device depends on the presence or absence of the single track table (**STT**) on the volume to be migrated.

If the STT is on the volume to be migrated, a cold start is recommended. This is necessary because the STT contains data which may be required for a number of JES3 functions (for example, the Volume Unavailable Table, DJC control blocks (dependent job control)). During a warm start, these records would be destroyed. Although the deletion of this volume would not necessarily be fatal, there is no guarantee that JES3 would be able to continue.

If the STT is not on the volume, the spool data set can be replaced over a warm start. You must decide:

- Whether to use a warm start or a 'warm start to replace a spool data set'.
- What to do about existing work on any volumes to be deleted.
- The sequence of actions to be undertaken.
- Will the new devices replace existing devices from the outset?
- What does that mean to existing work?
- Will the new devices be added to the complex?
- When will the old devices be deleted?

Options exist for handling the work present on a volume prior to the migration. You can drain the volume and allow the current work to complete or you can dump the volume to tape and restore the work after migration. You can delete jobs selectively prior to the migration to clean up the volume. Before deleting a volume, remember that JES3 will cancel all jobs in the system having spool data or allocation tables on the deleted volume.

If spool partitioning is used, spool data sets can be migrated selectively (that is, by partition). A spool partition is a logical grouping of spool data sets. Work which meets criteria may be directed to a subset of spool volumes or data sets.

Initialization statements like BUFFER, FORMAT, SPART, and TRACK contain parameters that can affect the performance of the JES3 spool configuration. Before migrating spool data sets, thoroughly understand how these parameters affect JES3 utilization of the devices used for spool data sets.

JES3 Data Sets

The JES3 Checkpoint Data Sets can be migrated over either a hot start or a warm start by taking advantage of the duplex checkpoint data set. During the start, one of the data sets can be replaced with the new device. If the duplex data set was not previously defined, it must be added during this start on the new device. JES3 will copy checkpoint records from the data set which was not replaced or added during the start to the data set which was replaced or added. The checkpoint data set which remains on the old device can be migrated to a new device during either another hot start or warm start, or, if duplexing is not desired, it can be deleted.

Note: Migration of JES3JCT to a new device requires a JES3 cold start.

For more information on the above topics and for help in determining spool space allocation requirements for the devices supported by the version and release of JES3 being used, see:

<i>MVS/Extended Architecture System Programming Library: JES3 Initialization and Tuning</i>	Shows how to determine size and placement of data sets including JES3JCT
<i>MVS/Extended Architecture Operations: JES3 Commands</i>	Describes how to add or replace checkpoint data set(s)

Other System Data Sets

New system dump data sets should be allocated. SYS1.LOGREC should be allocated and reinitialized. SMF's SYS1.MANn should be allocated and reformatted. VSAM's SYS1.STGINDEX must be re-defined.

The following system data sets contain load module output from the linkage editor. They can be allocated on the new device or copied by IEBCOPY or DFDSS. Both IEBCOPY and DFDSS utilize the COPYMOD function. COPYMOD may reblock data sets when they are moved to new device types.

SYS1.NUCLEUS
SYS1.CMDLIB
SYS1.DCMLIB
SYS1.IMAGELIB
SYS1.JES3LIB
SYS1.LINKLIB
SYS1.LPALIB
SYS1.SVCLIB
SYS1.TELCMLIB
SYS1.VTAMLIB.

Macro and parameter libraries are partitioned data sets, and can be moved using IEBCOPY or DFDSS.

SYS1.HELP
SYS1.INDMAC
SYS1.MACLIB
SYS1.PARMLIB
SYS1.PROCLIB
SYS1.SAMPLIB.

Catalogs

Do **not** move a catalog together with the data sets contained in it. This will avoid creating consistency problems that may occur if a failure happens during the move. Move user catalogs only in a controlled environment. There should be **no** other job, that access the catalogs being moved executing in the processor or processors.

Be certain to make backup copies of your catalogs before moving them.

With MVS/DFP Version 3, or a later release, and with MVS/XA DFP Version 2 Release 3, catalogs are immediately available when they are redefined. A locking mechanism is provided so that the system programmer can lock a catalog and examine or forward recover it before users have access to it. The RACF facility class IGG.CATLOCK allows the system programmer to gain access to a catalog and use it normally while it is locked.

More information on moving catalogs can be found in the appropriate *MVS Catalog Administration Guide* listed in the "Bibliography" on page 145.

Master Catalogs

A new catalog should be defined and the contents of the current master catalog may be copied using the access method services REPRO command. This is done as follows:

1. The target catalog is defined on the new volume.
2. The REPRO command is used to copy the master catalog's contents to the target catalog.
3. The reference to the master catalog volume is updated in PARMLIB.
4. The system IPL procedure is invoked.

Take this opportunity to convert your master catalogs to integrated catalog facility catalogs.

Integrated Catalog Facility Catalogs

Before moving any integrated catalog facility catalogs, make a backup copy using access method services or DFDSS. Integrated catalog facility catalogs can be moved with access method services or DFDSS.

VSAM Catalogs

VSAM user catalogs can be moved with access method services. Before moving a VSAM catalog, make a backup copy using access method services. Nonrecoverable VSAM catalogs can be moved with access method services REPRO. Recoverable VSAM catalogs can be moved with EXPORTRA or RESETCAT. Eventually, you should plan to convert your VSAM catalogs to integrated catalog facility catalogs.

OS CVOL Catalogs

OS CVOL catalogs can be moved using IEHMOVE or DFDSS. For a more detailed discussion on moving catalogs, see "Moving Catalogs with DFDSS" on page 89.

IMS/VS Data Sets

Use the IMS/VS utilities to move IMS/VS data sets. More information can be found in Chapter 9, “Moving IMS/VS and DB2 Data Sets” on page 103 and the IMS/VS manuals listed in the “Bibliography” on page 145.

DB2 Data sets

DB2 log data sets and the bootstrap data set (BSDS) can be moved with access method services or DFDSS. All other DB2 data must be moved by DB2 utilities or service facilities with the following exception: if you move all the data of a DB2 subsystem to a like device, then you can use DFDSS to move the entire volume containing the DB2 data. For more information, see Chapter 9, “Moving IMS/VS and DB2 Data Sets” on page 103 or the DB2 manuals listed in the “Bibliography” on page 145.

VSAM Application Data Sets

VSAM application data sets can be moved with access method services REPRO or EXPORT/IMPORT, or with DFDSS.

Non-VSAM Application Data Sets

Non-VSAM application data can be moved with the appropriate DFP utility or DFDSS. Some considerations for non-VSAM application data are:

- BDAM data sets can be moved with DFDSS or IEHMOVE; relative block and relative track integrity will be maintained. However, if the data set has the unmovable attribute, the moved data may not be usable by application programs containing device dependencies.
- ISAM data sets can be moved with IEBISAM or DFDSS. Now may be a good time to convert your ISAM data sets to VSAM.
- The space requirements for multivolume data sets should be examined to determine if they should remain on multiple volumes or if they should be allocated to a single volume.
- Data sets accessed through TSO are good choices for an early data move, because they are usually small, can be moved by DFDSS or by DFP utilities, and can be easily handled by volume or application groups. However, because the allocation and use of these data sets is not usually centralized, coordination with users is necessary to prevent negative results occurring with applications after moving the data sets.

Unmovable Data Sets

“Unmovable” data sets are so called because they were given the UNMOVABLE attribute when they were created. These data sets may need to be unloaded and reloaded by the application that created the data, although some unmovable data sets can be moved with DFDSS.

Data sets given the unmovable attribute are often accessed by an absolute record address generated by a program algorithm. In these cases, the program may need to be modified to operate correctly on the target devices when the source device has different characteristics.

Unmovable data sets have to be handled on an exception basis, and should not be processed while moving other data. If they must be located at a specific track and record address, it may be necessary to move them first, while that particular track

and record address are still available. It is preferable to consolidate them on volumes in a storage pool containing other unmovable data sets, handling them on an exception basis.

For more information, see "Moving Unmovable Data Sets with DFDSS" on page 90.

Data Sets Created by Non-IBM Programs

Before moving data sets written by non-IBM programs, determine whether that program supports the new device. Then, determine if IBM licensed programs can be used to move the data. If you are moving data between unlike devices, and you must use a non-IBM data movement program, verify that it supports unlike device types.

Tape Data Sets

Consider moving your small tape data sets to DASD storage. Some of the advantages of moving this type of data to DASD storage are:

- Reduced operator intervention for tape mounts
- Improved tape device utilization
- Better tape utilization
- Reduced access time
- Reduced job elapsed time
- System management of space, availability, performance and security.

For a more detailed discussion of moving many of the data types discussed in this chapter, see Chapter 8, "Moving Data in the MVS Environment" on page 87.

Considerations for the Placement of Data Sets

First decide where you will place your performance-critical data volumes. Then the placement and combination of data sets on your 3380 volumes can be planned, using the performance information gathered on the current configuration. Allow the system to determine the placement of data sets within a volume except where placement of a data set on a specific volume is important for I/O response or application dependent reasons. In that instance, you will have to manually place that data set on the volume. Using storage controls with cache will reduce the amount of hand-tuning required for data set placement.

Performance-Critical Data Sets

Pay attention to where the data sets were previously placed on highly active volumes. These data sets may have been placed near other data sets, such as the VTOC or catalog data sets, for performance reasons. Whenever possible, retain the performance advantages of previous tuning efforts.

Guidelines for the placement of performance-critical data sets include:

- Place concurrently active data sets on different paths to avoid contention for the same channel, storage control, and controller resources.
- Place concurrently active data sets on different volumes.
- Place small, concurrently active data sets on the same cylinder or adjacent cylinders when it is necessary to assign these data sets to the same volume (for example, the VTOC, the VTOC index, and the catalog).

- Allocate contiguous storage for a data set and then release the unused space to allow more closely placed active data sets.
- Avoid allocating on tracks for which alternate tracks have been assigned.
- Avoid shared DASD volumes except where sharing between systems is required.
- Reduce search activity in partitioned data set (PDS) directories by limiting the size of PDS directories.

Many of the performance guidelines just discussed involve trade-offs you must make for your specific applications.

With the MVS/DFP Storage Management Subsystem (SMS) the selection of devices based on performance requirements is automated for new data set allocations. Because SMS controls data placement, you are no longer required to perform the task of hand-placing sensitive data. Use the SMS Storage Class to specify a DASD I/O response time requirement and a caching bias of READ or WRITE I/O. This allows SMS to select the devices that can best meet the specified performance requirements for each data set.

Multivolume Data Sets

Find out why multivolume data sets are on multiple volumes. If a data set is on multiple volumes because of its size, you may be able to allocate the data set onto a single larger capacity 3380 volume. However, if the data set is spread across multiple volumes to maximize performance, you may choose to continue to spread the data set across multiple volumes.

Unmovable Data Sets

The unmovable attribute is used to indicate that a data set has location dependencies, either in the data set itself or in the code that processes the data set. Before moving an unmovable data set, you must determine where the dependencies are and whether or not the data set should be moved or re-created. In many cases, moving unmovable data sets will require programming changes.

If the unmovable data sets are interspersed on volumes that contain movable data sets, space management activities may suffer (for example, DFDSS DEFRAG). Consider consolidating the unmovable data sets on specific volumes or in a storage pool. In addition, if you assign your unmovable data sets in this way, they will not interfere with moving other volumes to 3380 volumes.

Planning a Data Management Strategy

In order to efficiently manage your storage subsystem, you need to review your current data configuration and determine requirements for each type of data set at your installation. The use of storage pools and storage groups will help you practice effective data management techniques and will also aid in device migration.

The following sections describe how pools and groups are implemented. *MVS SML: Managing Storage Pools* discusses using storage pools and storage groups in detail. An overview follows.

Using Storage Pools

Storage pools provide an effective way of managing data sets. A storage pool is a pre-defined set of DASD volumes used to store groups of related data. In a non-SMS environment, you refer to a storage pool using an esoteric name defined in the JCL. The number of pools you set up depends on your system requirements. The fewer pools that you have, the more efficient your storage management environment will be. Maintain as few pools as possible, so that automated storage management techniques can be applied to as much data as possible.

Storage pools will benefit device migration because they help maintain device independence and allow storage to be more centrally managed. Under SMS, storage pools eliminate the need for users to be aware of device types for their data. ACS routines direct the allocation of DASD volumes so that storage administrators will have to do less manual placement of data during the migration process.

Using Storage Groups

In an SMS environment, a storage pool is defined as a **storage group**. SMS provides valuable tools to set up and manage pooled storage. You can define management criteria for the different types of data in your installation using specific data classes. The values you specify in the classes identify your users' requirements for space, availability, and performance.

Using the Interactive Storage Management Facility (ISMF), you define four types of classes:

Data Class	A named list of data set allocation and space attributes that SMS assigns to a data set when it is created.
Management Class	A named list of management attributes that SMS uses to control DFHSM action for data set retention, migration, backup, and partial release.
Storage Class	A named list of data set storage service attributes that SMS uses to identify performance, availability and volume allocation requirements.
Storage Group	A named collection of candidate allocation volumes.

Automatic class selection (ACS) is the SMS mechanism used to assign SMS classes to new data sets, and to determine the storage group to which the data sets should be allocated. For example, if a data set has a requirement for continuous availability, you use storage class to specify this. This indicates to SMS that the data set should be allocated on a dual copy volume. Similarly, SMS can select a storage control that contains cache for new data set allocation, if the data set's storage class performance specification indicates this requirement and if cache is available.

In the past implementing storage pooling could require changes to coding standards for JCL. Implementing pooling for SMS-managed data requires no JCL changes. Because automatic class selection determines data placement, you are no longer dependent on user coded JCL parameters to efficiently manage storage.

Identifying Data Types not Supported by SMS

Some data types cannot be managed by SMS. If your data processing center uses any of these data types, you'll have to manage them separately. SMS does not support the following data types:

- Data sets cataloged in non-integrated catalog facility catalogs
- Data sets cataloged in more than one catalog
- Data sets that are not cataloged
- Non-integrated catalog facility catalogs
- CVOLS
- VSAM dataspace
- Unmovable data sets
- Data sets allocated with absolute tracks
- ISAM data sets
- Tape data sets.

A good long term plan to implement is to eliminate data sets that can't be managed by SMS. In the meantime, you should isolate these data sets from your SMS managed pools.

Documenting Configuration and Data Set Placement

Document your planned system configuration and data set placement for the following reasons:

- To enable others to review the planned configuration and data set placement. This helps prevent possible response time, throughput, or availability problems and helps make everyone aware of the priorities for different kinds of data on the system. Differences in viewpoint regarding the importance of different kinds of data may be resolved before moving the data.
- To record the system configuration and the general placement of data. This may be referenced during normal system operation.

Space Utilization Guidelines

For effective use of storage, your objective is to achieve a balance between space utilization and performance. Choices made regarding block size (physical record size) and data set placement affect this balance.

The amount of space on a track or cylinder actually occupied by user data depends to a large extent on the block size specified for the data set. Grouping logical records into blocks reduces the amount of space needed to store data. Since fewer physical records are needed to store the same number of logical records, the amount of space required for count and key areas, and for gaps between records is reduced. Blocking increases sequential processing efficiency by reducing the number of I/O operations needed to read or write the data.

For information on how to determine the number of fixed-length physical records that can be written on a 3380 track and cylinder, see *IBM 3380 Direct Access Storage Introduction*. It also describes a formula used to calculate space requirements, and includes an example of determining the amount of space to allocate for a data set.

Determining Space Requirements

Space requirements for data sets to be moved will be determined by:

- Logical record size
- Number of records
- Block size.

This discussion assumes that user data records are blocked and stored in equal-length physical records **without keys**.

If you are moving a data set from a volume on a DASD unit other than a 3380 volume, consider the different track capacities of these volumes. For example, a 3380 track can hold 47,476 bytes of customer data while a 3350 track can hold 19,069 bytes of customer data.

Next, determine the optimum block size and blocking factor for the application using the data. When moving data between unlike devices, reblock the data sets for optimum track utilization. If you are running under DFDSS Version 2 Release 3 or later, you can use the COPY command with the REBLOCK parameter to reblock your data sets to the most efficient block size.

Choosing an Efficient Block Size

Choosing efficient block sizes can improve space utilization and I/O performance. Block sizes must be chosen carefully to: meet the needs of the application, accommodate the size of the records being created, use space effectively, and guarantee device independence. Larger block sizes (6160 bytes and 9040 bytes) are more efficient in space utilization for a data set stored on a 3380 volume. The larger block size uses fewer tracks, and uses 3380 space more effectively because less space is taken up with track overhead (inter-record gaps and count areas).

The most efficient use of 3380 space (if you use IBM-supplied access methods) is obtained with a block size of 23,476 bytes. Two records of this size can be stored per track, with a space utilization of 98.9%. A block size larger than 23,476 means that only one record can be written per track. The remaining space on the track cannot be used to contain other data (assuming that all records in the data set are of equal length).

As block size decreases, the amount of data that can be stored on the track also decreases because smaller block sizes increase the amount of space needed for track overhead (inter-record gaps and count areas). For example, if you select a block size of 160 bytes (a blocking factor of two 80-byte logical records to one physical record), the 3380 track can contain 71 physical records for a total of 11,360 bytes of user data. In this case, the space utilization is 23.9%.

The selection of an optimum block size for a data set depends on:

- Track capacity
- Logical record size
- Data organization and accessing techniques
- Size and number of main storage buffers.

Select a block size appropriate for the track capacity of the device(s) you are using. Then adjust it to accommodate the size of logical records in the data set. For example, if the data set contains 80-byte logical records, a block size of 9040 bytes might be chosen. This results in a physical record that can contain 113 eighty-byte logical records. The 3380 track can contain five 9040-byte physical records, with a

space utilization of 95%. A block size of 900 bytes is a good choice for an all-3380 environment.

A block size of 6160 bytes is a good choice for physical records stored on various types of direct access devices at different times. A block size of 6160 bytes allows for two physical records per 3330 track (95.6%), for three physical records per 3350 track (98.9%), and for seven physical records per 3380 track (90.8%).

Using System-Determined Block Sizes with MVS/DFP Version 3

If you do not specify a block size or if you specify `BLKSIZE=0`, for SMS or non-SMS managed data sets, system-determined block sizes will be supplied for you. The data sets will be marked as reblockable so that, if you move data between unlike devices, the data sets will be reblocked automatically to the optimum blocksize for the new device.

Average Block Allocation: Using average block allocation rather than track and cylinder allocation allows for easy migration to new devices by making JCL and CLISTS device independent. To allocate in average blocks on the `SPACE` parameter, users specify the average number of bytes per record as the block length.⁴ The total number of records in the data set is specified as the number of blocks. The system then calculates the least number of tracks necessary to contain the number of blocks specified. Under MVS/DFP Version 3, you can also use the `AVGREC` subparameter to indicate the scale factor for the primary and secondary values. For example, an `AVGREC=K`, indicates that the primary and secondary allocation values are multiplied by 1024. So, to allocate enough space for a data set with 10,000 eighty-byte records the user specifies:

```
SPACE=(80,(10,1),RLSE),AVGREC=K
```

Figure 16. Using Average Record Length and the `AVGREC` Subparameter

Data Organization and Accessing Techniques

Small block sizes permit more concurrent operations on a channel, but they reduce the net data transfer rate (the actual amount of data transferred per second). Small block sizes may require more processor involvement for sequential data sets, and use more device space for a given amount of data.

Large block sizes allow a high net data transfer rate for sequential data sets and reduce the amount of processor time needed to process a channel program. A large block size requires exclusive use of the processor-to-DASD path during data transfer. However, larger data transfers can be justified, even for randomly accessed data sets, than were reasonable at the slower rate of earlier direct access devices.

Block sizes that result in satisfactory performance in a batch processing environment may not be desirable in an online, response-oriented application. In this case, it may be preferable to store less data per block in order to increase performance. The storage capacity of the 3380 helps to make this a reasonable compromise.

⁴ The average block value, however, does not necessarily mean the `LRECL` value.

The selection of block size may be affected by the size and number of main storage buffers used by your applications.

Developing a Data Movement Strategy

This section describes storage management and implementation practices that make moving data from one volume to another easier. Moving data to 3380 volumes does not require that an installation change the way it currently operates. You can successfully move data with the minimum level of software installed.

Minimize the number of simultaneous changes during data moves to limit the number of possible problems and to make any problems that do occur easier to solve.

If you are still using VSAM user catalogs and OS CVOLs, consider converting to integrated catalog facility catalogs on the 3380. In an SMS environment, VSAM catalogs are not supported, instead, data sets are cataloged using the integrated catalog facility. Once you convert to integrated catalog facility catalogs, move the data sets cataloged in them. This makes it easier to combine VSAM volumes on 3380 devices when the source volumes are owned by multiple VSAM user catalogs. For information on converting to integrated catalog facility catalogs, see the appropriate *MVS Catalog Administration Guide*.

Planning the Data Movement Sequence

Because of scheduling constraints and the amount of data being moved, it is unlikely that you can move the data in a single session. Therefore, you will need to devise a means of dividing your data into manageable groups, so that you can manage each group individually and, preferably, move each group in a single session.

If you use the DFDSS COPY command, it is not always necessary to schedule system time to move your data. Use the COPY command to move data sets or volumes of data during off-shift hours. DFDSS determines if a data set is in use. If it is, it won't be moved. This method will provide you with more time for data migration since you are not restricted to moving data during system down time.

Using Work Volumes

Before you move application data, consider using new devices as work volumes. These volumes should be mounted on devices associated with an appropriate esoteric name (for example, WORK) as defined in your sysgen. They should be mounted with the PUBLIC mount attribute. This process can be used to test and verify new devices.

SMS does not require that you define esoteric names for work data sets. Use the ACS routines to direct work data sets to the appropriate storage group that you've defined for the volume.

Plan to move the data that is the easiest to move and the least critical first. You may move more difficult and critical data when you gain experience in moving the data to the new devices.

Two Scenarios for Moving Data

Your plan for data movement should also depend on the sequence you use for installing new volumes and removing old ones. Two scenarios for carrying out these tasks are discussed here.

One scenario involves a gradual migration process in which a new volume is added to either a non-SMS pool or an SMS storage group, but old volumes remain. You can add a new volume either to an existing storage pool or group, or to a newly defined storage pool or group. No data movement is required at this time; the new volume will be selected for new data set allocation. The data sets will be balanced over the volumes in the storage pool or group.

When you want to remove a volume from a pool or group, begin by "draining" the volume. Draining a volume is the process of preventing new allocations and DFHSM automatic recalls on the expiring non-SMS and SMS volume(s) in the pool or group. Do this by mounting a non-SMS managed volume as private and specifying the DFHSM NOAUTORECALL, if the volume is a DFHSM primary volume. If the volume is SMS-managed, you can use SMS to quiesce allocations and DFHSM recalls.

At a later stage, it may be necessary to remove old volumes when they are no longer required in the storage pool or group.

Another scenario involves a one-for-one add and replacement procedure of volumes in a storage group or pool. This means that the installation of a new volume depends on the removal of an old volume from the storage pool or group. The data that was on the old volume must be moved to the new volume.

Once you have chosen a scenario for managing volumes, you can consider the types of data at your installation and the movement sequence that would be most appropriate to move them.

Moving Specific Data Types

When grouping the data sets, take into account the kind of data you have to move, scheduling constraints, the potential impact of moving that collection of data and any other considerations unique to your system. The following are some possible ways to move the data:

By Volume

Moving entire volumes should be done when the volume contains data that requires no reorganization or preferential treatment for performance (or other) reasons. Generally, moving entire volumes may be the quickest way to move data.

By Application

Moving all data for a given application simultaneously is especially appropriate for an application that has already established data set naming conventions. Naming conventions permit the use of the filtering capabilities of DFSS to select all data sets of a particular application. Inform the owner of the application data of the data migration plans.

By Catalog

If data sets are logically grouped in separate catalogs, consider moving these groups together. This method is especially appropriate for moving VSAM data sets currently residing on multiple volumes.

By Individual Data Set

Moving data sets individually is usually required when you need specific placement for performance-oriented data or you have to move data that needs special handling (such as “unmovable” data sets, or multivolume data sets), and system data sets.

By Data Base

Online data base data sets are often very large and have high availability requirements. They should be moved during non-production time to limit the impact on your environment. Online data bases often have established procedures in place for their backup and recovery. Use these existing procedures to move the data bases.

Alternate Methods to Put Data on 3380s

There are ways to add devices to your system and fill them with data without copying data sets directly to them. You may allocate new data or extensions of existing data to the new devices. With generation data groups or chained transaction files, you can move from old to new devices by cataloging the new generations or data sets on the new devices, and allowing the old devices to be freed up during the generation cycle. In an SMS environment, you can use ACS routines to redirect data sets to new volumes. New volumes can be defined to existing storage groups or new storage groups.

Data sets can be moved to the new devices during normal processing by allocating new data sets to the new devices. Change the JCL to direct the new data sets to the new devices. This will move those data sets during scheduled production runs, gradually filling the new devices; leaving only the residual data sets to be handled at the end of a processing cycle.

If you are working in an SMS environment, there is no need to change JCL to redirect data sets. Data sets are redirected via ACS routines.

Moving Data Sets with DFHSM

DFHSM can be used by DFHSM-authorized users to move data sets between like and unlike devices. DFHSM stores data sets in a device-independent format during migration and can move data sets as part of the normal migrate/recall process.

In addition, you can use the DFHSM MIGRATE command with a CONVERT parameter to simplify moving data to a different level 0 user volume. The CONVERT parameter causes DFHSM to migrate and then immediately recall a specified data set (or all data sets) on a volume to another volume or group of volumes on like or unlike devices.

If you are operating under SMS, ACS routines will be invoked for all data sets to determine whether a data set is to be recalled to SMS managed storage or not. If a data set is to be recalled to SMS storage, SMS will determine the volume to which the data set will be recalled, regardless of whether a target volume was specified on the MIGRATE command. The target volume is only used as a filter input to the ACS routines.

To add an SMS managed volume to the storage hierarchy managed by DFHSM, you must define it to a storage group using ISMF.

Use the ADDVOL command to add a new non-SMS managed volume to the hierarchy. DFHSM will then use this volume for the DFHSM storage management

function for which it was defined. For example, you can define a new volume as a primary volume that is a candidate for recalled data sets.

You can also use the ADDVOL command with the DRAIN parameter to empty a DFHSM migration volume over a period of time. DFHSM can still recall data sets from the volume, but can not migrate data sets to that volume.

There are two ways to use the FREEVOL command to empty DFHSM DASD migration volumes. Specifying FREEVOL causes DFHSM to move all the data sets on a migration level 1 volume to other migration level 1 volumes or to migration level 2 volumes. You can also use this command to empty mass storage volumes that are defined as DFHSM migration volumes.

For details on what data sets are supported by DFHSM and how to move data sets with DFHSM, see the appropriate DFHSM manuals listed in the "Bibliography" on page 145.

Planning for Backup and Recovery While Moving Data

It is important to plan and implement a strategy for backup and recovery while moving the data.

Even though you have backup and recovery procedures tested and in place, you will probably have to set up additional procedures to protect the installation's data during the data move, when it is especially vulnerable.

The following are suggested:

- Back up each old volume before its contents are moved.
- Back up each new volume as soon as its contents are complete.
- Keep written records that will enable you to quickly tell where any data set on any moved volume has been relocated, and where its original backup (from the old volume) and its latest backup are located.

You can use DFDSS to generate backup copies of volumes when moving volumes, or to generate backup copies of individual data sets when moving data sets. Use data base utilities to back up IMS, DB2, and CICS data bases.

For examples of using DFDSS to create backup tapes and to restore volumes or data sets, see "Backup and Recovery While Moving Data" on page 88

Note: The same version of DFDSS must be used for backup and recovery.

Documenting and Reviewing Your Data Movement Strategy

After you have determined your data movement strategy, document it. A checklist of installation and data movement planning activities is contained in Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121. When you have completed the activities that pertain to your installation, have the data movement plan reviewed. Include the following in your documentation:

- Diagrams describing how the new devices will fit into your system configuration and what each volume's prime use will be (for example, system volume).

- A description of where the data sets will be placed (preferably in terms of storage pools). Where needed for performance reasons, state on which volume the specific data sets are to be placed.
- A description of the data groups you have selected to move.
- The data movement sequence.
- A description of the supplemental backup and recovery procedures that will be used.

When you have completed the documentation, have it reviewed by:

- Management. Management will normally be involved in the review and approval process. They should provide sufficient system resources needed for moving the data.
- User groups affected by the move plan. Problems caused by data being unavailable while moving the data can be avoided by a review of the data movement plan. Also, informed users are generally more cooperative with the data move activities than uninformed users.
- The data move team. Use the document to coordinate the team's efforts. With a written plan in place, simultaneous activity by several team members is possible.

By documenting and reviewing the data move plan, you can handle potential problems before, rather than after, the data move effort.

Scheduling System Time

After your data move plan has been approved and you have installed the new 3380 units, schedule an appropriate amount of system time to move the data. You should now have a reasonably accurate idea of how much time will be needed to move each group of data, and when the time will be needed.

For more information, see the MVS Storage Management Library in the "Bibliography" on page 145.

Chapter 6. Installing the 3380

Having completed the planning activity described in Chapters 2 through 5, you're ready to install the 3380 units on your MVS system. The DASD installation process includes the following tasks:

- Defining the 3380 units to the system
- Physically installing the hardware
- Preparing the 3380 volumes for use

This chapter provides information to help you complete these tasks.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a DASD installation section that you can use to track the completion of these activities. When attaching 3380s to a 3990 storage control, you can use the configuration worksheets in *IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide*.

Defining the 3380 to the System

Before you can use the 3380, you must identify its existence and location to MVS. Do this before you actually install the 3380 units, so that you can test the configuration definition.

Defining the 3380 to the operating system includes three steps:

1. Assigning I/O addresses or device numbers to the new units
2. Defining the units to MVS (by running MVS configuration program or performing an iogen)
3. Defining the units to the processor (by running the I/O Configuration Program, if necessary)

The following sections describe these steps in more detail. For basic information about the system generation process, see the appropriate MVS/XA or MVS/370 *Installation: System Generation* reference manual, or the *MVS/ESA System Generation*.

Assigning I/O Addresses or Device Numbers

MVS/370 systems use **I/O addresses** to identify each device (volume) in the 3380 unit. MVS/ESA and MVS/XA systems use **device numbers** to accomplish the same purpose.

You will need to assign I/O addresses or device numbers for all the devices in your new 3380 units and supply these addresses or numbers to the system. You will also need to provide these values to your service representative so the 3880, 3990, and 3380 units can be set to recognize them.

If you plan to install more 3380 units in the future, you may want to assign their addresses or numbers now. You can define these units to MVS during the configuration definition process, and at IPL time MVS will mark the units that have not been installed as offline.

MVS/370 uses the I/O address to identify both a 3380 device and the path used to reach it; thus, a given 3380 device can have several different addresses, get for each path by which it can be reached.

MVS/ESA and MVS/XA use the device number to identify a 3380 device only; because the path is selected dynamically by the channel subsystem, there is no specific path identifier within the device number.

Detailed information on how to assign I/O addresses or device numbers for the 3380 is discussed in *IBM 3380 Direct Access Storage Introduction*. Corresponding information for the 3380 Model CJ2 is found in *IBM 3380 Direct Access Storage Direct Channel Attach Model CJ2 Introduction and Reference*.

Defining the I/O Configuration to MVS

To add 3380 devices, you:

- Define the new devices to MVS and to the channel subsystem
- Add the devices to the appropriate storage pool if using system-managed storage
- Assign mount and use attributes

Defining New Devices to MVS and the Channel Subsystem

You define storage subsystems to the channel subsystem using CNTLUNIT statements for the storage control and IODEVICE statements for the 3380 with the IOCP program. In addition, the 3380s are defined to MVS using these same IODEVICE statements with the MVS configuration program available with MVS/SP 2.2 (or with the sysgen process in prior releases of MVS/SP). Since a 3380 Model CJ2 contains storage control function and DASD, the CJ2 is defined with both CNTLUNIT and IODEVICE statements.

You may create a common input stream for both the IOCP program and the MVS Configuration Program (or sysgen) by merging the CNTLUNIT and IODEVICE statements. See the *MVS/XA or MVS/ESA MVS Configuration Program Guide and Reference* for further information.

IODEVICE Statements for 3380 Devices: Adding 3380 devices to your system requires coding IODEVICE statements which you will use as input to IOCP and either iogen for MVS/370 or the MVS configuration program for MVS/ESA and MVS/XA. For the IODEVICE UNIT parameter, specify 'UNIT=3380' for all 3380 models.

If you already have UCBs generated for 3380 Model AA4 units, you can substitute Model AD4 units without performing an iogen. Similarly, if you have UCBs generated for 3380 Model AE4 or BE4 units, you can substitute AJ4, BJ4, AK4, BK4, or CJ2 units without performing an iogen. Updating the existing I/O configuration using MVSCP (available with MVS/SP 2.2 or later releases) or performing an iogen with the operating system at the level supporting these models is required when the size of the UCB is changed, but not when only the contents of the UCB are changed. Failure to do this will cause unpredictable results.

Installing 3380 AJ4 or AK4 strings attached to 3990 Storage Controls in the MVS/XA DFP 1.1.3 and MVS/370 DFP 1.1.2 environments requires an iogen as part of the installation process, but an iogen is not required in the MVS/XA DFP 2.2.3, MVS/XA DFP 2.3.0 or a later release, and MVS/DFP 3.1.0 or later (except as noted previously).

This section provides supplementary information in the form of examples. For detailed information, see the appropriate MVS/370 or MVS/XA *Installation: System Generation* or the MVS/ESA *System Generation and IOCP User's Guide and Reference*.

IBM 3380 devices can be installed in four MVS environments:

- MVS/370 operating system with 303x, 4341, 4361, 158, 168, or 9370 processor⁵
- MVS/370 operating system with 370 Extended Architecture processor (308x, 3090 or 4381)
- MVS/XA operating system with 370 Extended Architecture processor (308x, 3090 or 4381)
- MVS/ESA operating system with ESA/370 processor (3090E or 4381E)

The processors and the devices and storage controls which attach to them are shown in Figure 3 on page 3.

To add the four 3380 units represented in Figure 17 to processors in each of the MVS environments listed above, use the statements shown in Figure 19 through Figure 21.

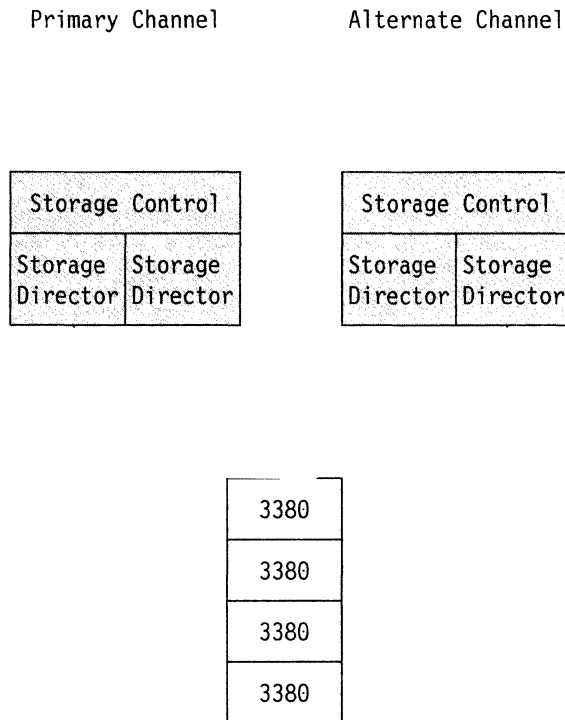


Figure 17. Configuration Assumed for MVS System Generation Examples

Figure 3 on page 3 summarizes the attachment possibilities of 3380 models to storage controls to IBM processors. For information on how to configure the different 3380 models, see Chapter 4, "Planning the Hardware Configuration" on page 29 and *IBM 3380 Direct Access Storage Introduction*.

⁵ Non-XA operating systems can be run with the 3090 models 300E, 500E, and 600E using the PR/SM feature.

Note: It is important to note that the CNTLUNIT statements must match the physical storage controls on a channel on a one-for-one basis, and must define all units that the service representative has set the storage director (3880) or storage path (3990) to recognize. If you do not define that full range of addresses, or if you have more than one CNTLUNIT statement per storage control, the result can be missing interrupts (which can cause device-end-pending problems).

MVS/370 Operating System with 370 Architecture Processor: This environment requires only an iogen, no IOCP generation is needed. The iogen statement in Figure 18 shows how to define 3380 volumes.

```
DASD IODEVICE UNIT=3380,ADDRESS=(D60,16),OPTCHAN=F,
      FEATURE=ALTCTRL
```

Figure 18. Sample 3380 IODEVICE Statement

MVS/370 Operating System with 370 Extended Architecture Processor: In addition to an iogen with the UNIT parameter of '3380', an IOCP generation is necessary with 308x and 3090 Extended Architecture processors. This operation is required to load the input/output configuration data set (IOCDs). The OPTCHAN parameter is used to specify alternate paths.

The IODEVICE statement shown in Figure 18 shows how to define sixteen 3380 volumes on channel D, with channel F as an alternate channel.

To define the devices in Figure 18 to IOCP, use the statements in Figure 19. The CHPID statement associates paths 15, 16, and 17 with channels D, E, and F respectively. The third subparameter of CHPID defines the channel set. The CNTLUNIT and IODEVICE statements define path 15 and path 17 through two 3880 storage directors (both addressed as 6) to 3380 unit addresses D60 through D6F and F60 through F6F respectively. Data streaming is specified by PROTOCOL=S and is only valid for a storage control attached to a block multiplexer channel. SHARED=N specifies that the storage control supports multiple I/O requests concurrently (one for each attached device),

```
CHPID=((15,D,0),(16,E,0),(17,F,0),...)
CNTLUNIT CUNUMBER=0D6,UNIT=3880,UNITADD((60,16)),
        SHARED=N,PROTOCL=S,PATH=(15)
CNTLUNIT CUNUMBER=0F9,UNIT=3880,UNITADD((60,16)),
        SHARED=N,PROTOCL=S,PATH=(17)
IODEVICE UNIT=3380,ADDRESS=(D60,16),CUNUMBER=(0D6,0F9),
        OPTCHAN=F
```

Figure 19. IOCP Generation, MVS/370 Operating System with 370 Extended Architecture Processor. Ellipses (...) indicate other channel paths that are not shown.

MVS/XA or MVS/ESA Operating System: To add 3380s to either an MVS/ESA system on an ESA/370 processor or an MVS/XA system on a 370/XA processor, perform an IOCP generation to define the DASD to the processor. Note the OPTCHAN parameter is not required in these environments. Use MVSCP or an iogen to update the I/O configuration to the operating system.

Figure 20 on page 69 shows the IODEVICE statement required to define sixteen 3380 volumes with addresses 6A0 through 6AF, which will have alternate paths and

the ability to be shared with other systems and channels. This statement can be added to the input streams to both IOCP and MVSCP.

```
DASD IODEVICE UNIT=3380,ADDRESS=(6A0,16),  
FEATURE=(ALTCTRL,SHARED)6
```

Figure 20. IODEVICE Statement for MVS/XA or MVS/ESA

CHPID, CNTLUNIT, and IODEVICE statements defining path 06 and path 07, through two 3880 storage directors addressed as A, to 3380 unit addresses A0 through AF, are shown in Figure 21. The 16 volumes will be generated as addresses 6A0 through 6AF.

```
CHPID=((06,6,0),(07,7,0),...)  
CNTLUNIT CUNUMBER=060,UNIT=3880,UNITADD=((A0,16)),  
SHARED=N,PROTOCL=S,PATH=(06)  
CNTLUNIT CUNUMBER=070,UNIT=3880,UNITADD=((A0,16)),  
SHARED=N,PROTOCL=S,PATH=(07)  
IODEVICE UNIT=3380,ADDRESS=(6A0,16),CUNUMBER=(070,060)
```

Figure 21. IOCP Generation, MVS/XA or MVS/ESA. Ellipses (...) indicate other channel paths that are not shown.

Data streaming is specified by PROTOCOL=S and is only valid for a storage control attached to a block multiplexer channel. SHARED=N specifies that the storage control supports multiple I/O requests concurrently (one for each attached device),



Four-Path Connection to MVS/XA or MVS/ESA

To add 3380s to either an MVS/ESA system on an ESA/370 processor or an MVS/XA system on a 370/XA processor, perform an IOCP generation to define the DASD to the processor. An OPTCHAN is not required in these environments. Use MVSCP to update the I/O configuration to the operating system. Include all paths available in the PATH parameter of the CNTLUNIT statement and include all addresses for the 3380 in the string in the ADDRESS parameter of the IODEVICE statement.

Figure 22 shows the configuration that is defined by the statements in Figure 23.

⁶ The ALTCTRL parameter is not required by MVS/XA or MVS/ESA, but can be included in the IODEVICE statement.

Channels Channels



MPSD = Multi-path
Storage Director
SP0-3 = Storage Paths 0-3
A1-A4 = Controllers 1-4

Figure 22. Four-Path Configuration Example

Figure 23 shows the statements that define channel path ids 1, 2, 11, and 12, attaching a 3990 storage control and thirty-two 3380 volumes. The 32 volumes will be generated as addresses 180 through 19F. An additional 32 addresses have been reserved for a future non-disruptive device install. Note that the UNIT parameter of the CNTLUNIT statement is coded UNIT = 3990, as this configuration includes a 3990 Storage Control.

```

CHAN1 CHPID PATH=((01,1,0),(11,1,1)),TYPE=BL
CHAN2 CHPID PATH=((02,2,0),(12,2,1)),TYPE=BL
CNTL180 CNTLUNIT CUNUMBR=180,UNIT=3990,UNITADD=((80,64)),
        PATH=(01),PROTOCL=S,SHARED=N
CNTL280 CNTLUNIT CUNUMBR=280,UNIT=3990,UNITADD=((80,64)),
        PATH=(02),PROTOCL=S,SHARED=N
DDEV180 IODEVICE UNIT=3380,ADDRESS=(180,64),
        FEATURE=(SHARED,ALTCTRL),CUNUMBR=(180,280)

```

Figure 23. IOCP Generation for a Four-Path 3380-3990 Configuration

Data streaming is specified by PROTOCOL = S and is only valid for a storage control attached to a block multiplexer channel. SHARED = N specifies that the storage control supports multiple I/O requests concurrently (one for each attached device),

Defining a 3380 Model CJ2 String to an Extended Architecture Processor

A 3380 Model CJ2 attaches directly to a processor channel. A CJ2 string is defined to the processor as if it were a 3990 Model 1 and a 3380 string with 14 devices starting with device 2. Figure 24 shows the statements that define a 3380 Model CJ2 string, as shown in Figure 15 on page 46, attaching directly to channels 15 and 16.

The CHPID statement associates paths 15 and 16 with channels D and E respectively. The third subparameter of CHPID defines the channel set. The CNTLUNIT and IODEVICE statements define path 15 and path 16 through the two

storage paths in the C-unit to 3380 unit addresses D62 through D6F and F62 through F6F respectively.

Note: Although a full 3380 string headed by a CJ2 has only 14 devices, the string length is defined as if 16 devices were there.

```
CHPID=((15,D,0),(16,E,0),...)
CNTLUNIT CUNUMBER=0D6,UNIT=3880,UNITADD((60,16)),
        SHARED=N,PROTOCL=S,PATH=(15)
CNTLUNIT CUNUMBER=0F6,UNIT=3880,UNITADD((60,16)),
        SHARED=N,PROTOCL=S,PATH=(16)
IODEVICE UNIT=3380,ADDRESS=(D60,16),CUNUMBER=(0D6,0F6),
        OPTCHAN=E,FEATURE=ALTCTRL
```

Figure 24. IOCP Generation, 3380 CJ2 String. Ellipses (...) indicate other channel paths that are not shown.

Data streaming is specified by PROTOCOL = S and is only valid for a storage control attached to a block multiplexer channel. SHARED = N specifies that the storage control supports multiple I/O requests concurrently (one for each attached device),

Defining a 3380 Four-path, Two-Path Intermixed String

This section shows how to define an intermixed configuration containing:

- A 4-path string of Enhanced Subsystem 3380s (J-units and K-units)
- Two 2-path strings of Extended Capability 3380s (D-units and E-units)
- 3990 Model 2 or 3 Storage Control.

With the configuration shown in Figure 25 on page 72, any channel path connected to the 3990 Storage Control can transfer data to or from any of the attached 3380 strings.

as many as 8 channels as many as 8 channels

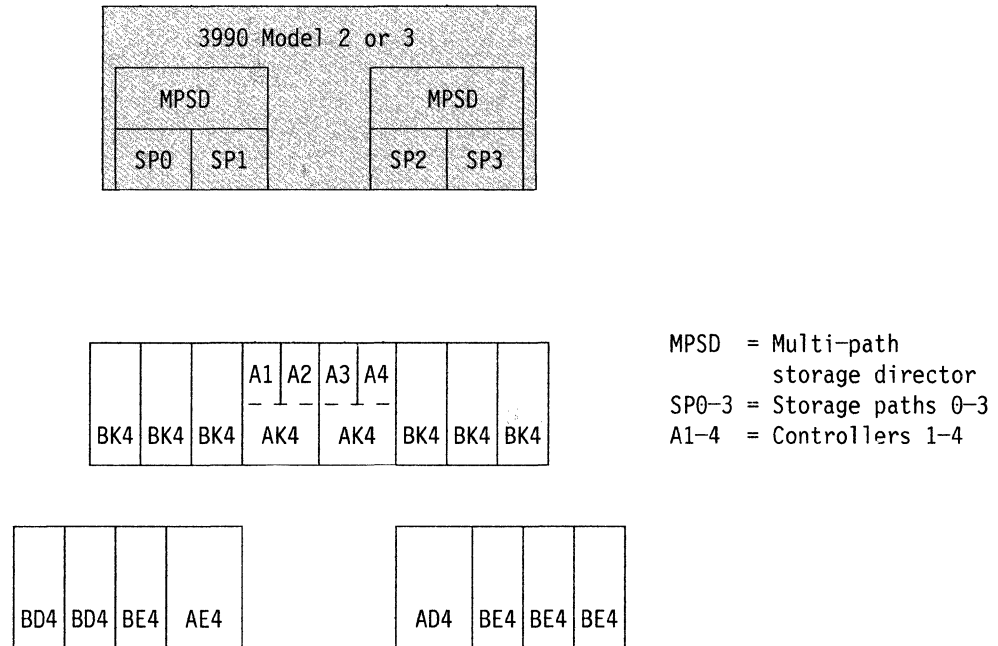


Figure 25. Four-Path, Two-Path Intermixed String Configuration Example

Defining a 3380 Four-path, Two-Path Intermixed String with IOCP

Figure 26 shows how to define the configuration shown in Figure 25 IOCP and either iogen for MVS/370 or the MVS configuration program for MVS/XA or MVS/ESA. Figure 26 could be defined in the IOCP source file.

```

CNTLUNIT CUNUMBR=128,UNIT=3990,UNITADD=((80,64)),SHARED=N,
        PATH=28,PROTOCL=S
CNTLUNIT CUNUMBR=12A,UNIT=3990,UNITADD=((80,64)),SHARED=N,
        PATH=2A,PROTOCL=S
IODEVICE CUNUMBR=(128,12A),UNIT=3380,ADDRESS=(1880,64)
    
```

Figure 26. IOCP Generation for Four-Path, Two-Path Intermixed String

Storage Pools

Storage pools are set up with MVSCP available with MVS/SP 2.2 (or the sysgen process in prior releases of MVS/SP). The UNITNAME statement defines and identifies the storage pools. In addition to adding the devices you are installing now to your storage pools, if you intend to install more 3380 units later and have specified them in the input stream to MVSCP, you can also add these anticipated devices to storage pools at this time. More detailed information on sysgen process can be found in the appropriate MVS system generation manual listed in the "Bibliography" on page 145.

Mount and Use Attributes

The “mount” and “use” attributes of a DASD volume are defined to MVS in the SYS1.PARMLIB member VATLSTxx. Define each volume with a use attribute of PUBLIC, STORAGE, or PRIVATE. Storage pools defined during the MVSCP or sysgen process can be subdivided by using these attributes.

Defining the 3380 to SMS

Use MVSCP with IODEF mode to define the full range of device addresses to be used by the system. It is not necessary to specify the UNITNAME information for SMS storage pools at sysgen time. Furthermore, if the volumes are SMS-managed, there is no need to define storage attributes in the VATLSTxx member.

Adding an SMS volume: All volumes in an SMS storage group must be of a like device type. When adding a new SMS volume, you may need to define a new storage group using the ISMF Storage Group Application. If you are defining a new storage group, you may want to consider pre-defining a list of volsers for the storage group to ease installation of volumes in the future.

To add volumes to a storage group, invoke ISMF to do the following:

1. Update the source for the currently active configuration to define the volume to the storage group
2. Activate the updated configuration
3. Initialize the volume via the Storage Group Application.

If the volume has been pre-defined to the storage group, you can bypass steps 1 and 2 and initialize it using the Storage Group Application or ICKDSF.

More information on storage pools and activating the SMS configuration can be found in the *MVS SML: Managing Storage Pools* and the *MVS/ESA Storage Administration Manual* respectively.

Installing the 3380 Units

After defining the 3380s to the system, you can physically install them. For information about cabling, space, weight, environmental, and other physical considerations, see *IBM Input/Output Equipment Installation Manual—Physical Planning*.

Preparing the Volumes for Use

After the service representative physically installs the 3380s, you must prepare the volumes for use by MVS. The 3380 is delivered without a volume label or volume table of contents (VTOC), which you supply during device initialization. The following sections describe these activities.

Calculating the Size of VTOC and VTOC Index

When installing a new 3380, you determine the size of the VTOC and the VTOC index. This information is needed by Device Support Facilities (ICKDSF) during initialization.

The VTOC is composed of data set control blocks (DSCBs). A 3380 cylinder can contain 795 DSCBs (53 per track times 15 tracks). Because of the increased capacity of 3380 volumes, the size of the VTOC deserves special consideration.

The 3380 Models AE4, BE4, AK4, and BK4, can hold more data sets and may require larger VTOCs and VTOC indexes than lesser capacity 3380 volumes.

Note: Because there are fewer total tracks on 3380 single capacity volumes than on 3350 volumes, a 3350 volume filled with single-track data sets cannot fit on a standard 3380 volume (although only the first 40% of each 3380 track will be used). If you plan to allocate more data sets on a volume, verify that the sizes of the VTOC and VTOC index are adequate for the planned usage. For more information, see *MVS Storage Management Library: Managing Storage Pools*.

Calculate VTOC Size

The size of the VTOC for 3380 volumes can be calculated as follows:

1. Estimate the maximum number of data sets that will reside on the volume.
2. Divide that number by 53, and round up, to find the number of tracks necessary for the VTOC.

Calculate VTOC Index Size

The VTOC index is a separate data set containing four record types that have information on volume and VTOC status. Systems that use indexed VTOC support can realize performance improvements in the areas of volume space management and input/output operations to the VTOC. Because of performance gains, the use of indexed VTOCs on all your volumes is strongly recommended.

The calculation of the size of the VTOC index depends on the average length of the data set names expected for the volume, and the maximum number of data sets expected for the volume. If you do not want to calculate the size of the VTOC index, you can omit the third subparameter of the INDEX parameter of the INIT statement, and ICKDSF will calculate an index size for you. However, before choosing to do this, see *Device Support Facilities User's Guide and Reference* for the assumptions on which the VTOC index size calculation is based and to find out how to precisely calculate the size of the VTOC index.

Initializing New Volumes

After calculating the VTOC and VTOC index sizes, initialize the volumes with ICKDSF. Device Support Facilities Release 10.0 provides two levels of initialization for 3380s:

Minimal level creates the volume label, VTOC, VTOC index, and, optionally, creates IPL records. Specifying the NOVALIDATE parameter in the INIT command causes minimal initialization.

Medial level performs all functions of minimal level, plus the validation of each track's home address (HA) and record zero (R0). Specifying the VALIDATE parameter in the INIT command causes this volume verification. Volume verification can also be performed by the INSTALL and REVALIDATE commands (which include other ICKDSF functions). If INSTALL or REVALIDATE are used, they must be followed by a minimal initialization.

The devices are shipped from the plant of manufacture with home address (HA) and record 0 (R0) formatted on all tracks, and no alternate tracks assigned. When the unit is installed, the service representative tests the unit to ensure that data

can be read and written. For guidelines on how to initialize 3380 volumes, see *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*.

When the 3380 arrives from the factory, there are no volume labels on the devices. Use ICKDSF to initialize the device offline⁷, label the volume, and write the VTOC and VTOC index.

When initializing a volume with a VTOC index, place the VTOC index adjacent to the VTOC for best performance. You may calculate and specify the number of tracks, or you can omit that subparameter and ICKDSF will calculate the size of the index.

Initializing Volume Offline

An offline initialization at the minimal level of a new 3380 unlabeled volume (where approximately 2000 data sets will be placed) is shown in Figure 27. The UNIT(723) parameter specifies the device address or device number of the device to be initialized and is required when initializing a volume offline. The NVFY parameter causes the verification of volume serial number and owner identification to be bypassed.

The VTOC is defined on cylinder 435 starting at track 4, and continues onto cylinders 436, 437, and 438. The VTOC index, adjacent to the VTOC, will reside on cylinder 435, track 0, and is 4 tracks in size.

```
//STEP01 EXEC PGM=ICKDSF,REGION=512K
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
INIT UNIT(723) NVFY PURGE VOLID(338001)
VTOC(435,4,56) INDEX(435,0,4)
/*
```

Figure 27. Offline Initialization of 3380 Volume

⁷ If you receive error message ICK318511, VARY each device online (ignore the system message in response to this VARY). This conditions the operating system and channels to the device characteristics but the device remains offline.

Chapter 7. Operating the 3380 under MVS

This chapter contains information an operator needs to know to use the 3380 under MVS. You can use the information in this chapter to prepare procedures that can be used by your system operators.

Most changes to the operational status of the 3380 are controlled from the system operator console. System commands are issued from the system operator console to control 3380 status and to determine the status of the 3380s in your storage subsystem.

Some 3380 operation status changes are controlled from the operator control panel. A control panel for each 3380 string is located on the end cover of 3380 Model A04, AA4, AD4, and AE4 units; or on the front cover of 3380 Model AJ4, and AK4 units. B-units do not have operator control panels. The 3380 Model CJ2 operator panel is located on the front cover of the unit. The operator control panels are described and illustrated in *IBM 3380 Direct Access Storage Introduction* where you will also find information on:

- Reading the operator control panels
- Turning the power on and off
- Enabling and disabling controllers and devices.

This chapter describes the following operator tasks for the 3380 as they are performed in the MVS environment:

- Displaying 3380 status
- Varying volumes online and offline.

In an SMS environment, the VARY and DISPLAY commands can be used to monitor and change status of SMS volumes. The DEVSERV command can be used to determine storage group information for a device. For more information on how these commands are used in an SMS environment, see *MVS/ESA System Commands* manual.

Appendix A, "Sample DASD Installation and Migration Project Plan" on page 121 contains a checklist, Planning Operations Changes, you can use to track the completion of activities described here.

Using the DISPLAY Command to Determine Device Status

You can use the DISPLAY command to determine the state of a 3380 device, string controller, 3880 storage director, 3990 storage director, or channel. The DISPLAY command, when issued in an MVS/ESA and MVS/XA operating system environment, provides you with more information than when issued in the MVS/370 operating system environment. This additional information will give you a better understanding of the state of the DASD subsystem.

The DISPLAY command syntax is described in the appropriate *System Commands* manual for your MVS operating system. This section provides examples of what a system operator might issue, what the operator could expect to see as a response, and how to interpret the response.

The DISPLAY U Command

You use the DISPLAY U command to determine the status of a device. During normal operations, device status changes as it responds to channel commands. The DISPLAY U command may or may not always return status change information during normal operation. An indication that the device state might not be changing is a system console message:

In MVS/370 systems, message:

IGF990I DEVICE END MISSING

In MVS/ESA and MVS/XA systems, message:

IOS071E START PENDING

When you issue the DISPLAY U command to a DASD, message IEE450I is returned. The message contains the status of the device and can help the operator determine whether an address (a device) is working properly or not. The possible status conditions include online, offline, allocated, busy, and not ready. Other status conditions indicate whether the device is reserved and whether the volume contains system resident data sets.

When you issue several DISPLAY U commands and you consistently receive a START PENDING message, this may indicate an out-of-sync condition (this condition is described in "Out-of-Sync Conditions" on page 84).

However, if you are operating in a shared DASD environment, START PENDING status may indicate that a job running on another system has caused the busy condition. For example, another system may be performing a DFDSS DUMP/RESTORE or link-edit routine. In this case, the DISPLAY U command would reveal an A-BSY-R or O-BSY-R status for the system that caused the busy condition. If the requesting system receives a START PENDING status and there is no A-BSY-R or O-BSY-R on any other system, then it is possible that an out-of-sync condition exists.

Figure 28 shows you the result of issuing a DISPLAY U command for devices 281 through 284.

```
DISPLAY U,,,281,4
IEE450I 10.04.44 UNIT STATUS 385
515 UNIT TYPE STATUS VOLSER VOLSTATE UNIT TYPE STATUS VOLSER VOLSTATE
515 281 3380 0 338001 STRG/RSDNT 282 3380 0 338JT1 PUB/RSDNT
515 283 3380 0 338JT2 PUB/RSDNT 284 3380 0 338JT3 PUB/RSDNT
```

Figure 28. Example of DISPLAY U Command

In this example, devices 281-284 are online. If one or more devices had a status of A-BSY-R, repeat the DISPLAY U,,,281,4 command. If any device continues to show an A-BSY-R status, assume that an out-of-sync condition exists. If the system operator cannot correct this out-of-sync condition by issuing 'VARY 281, ONLINE', call your systems programmer. If the VARY command corrects the out-of-sync condition, vary online the other devices attached to the same storage director of a 3880 or storage path of a 3990.

The DISPLAY M Command in MVS/XA and MVS/ESA

Use the DISPLAY M command in the MVS/XA and MVS/ESA environment to list information about all paths to a set of devices or to an individual device.

The response to a DISPLAY M=CHP(nn) command issued in an MVS/XA or MVS/ESA operating system is shown in Figure 29. Information for all devices on channel path 43 is displayed. Symbols are explained at the bottom of the display.

```

080 DISPLAY M=CHP(43)
080 IEE174I 12.42.47 DISPLAY M 742
000 CHANNEL PATH 43 STATUS BY DEVICE
000  0 1 2 3 4 5 6 7 8 9 A B C D E F
000 30 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 31 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 32 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 33 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 34 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 35 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 36 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 37 $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 38 + + + + + + + + + + + + + + +
000 39 + + + + + + + + + + + + + + +
000 3A $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 3B $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 3C $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 3D $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 3E $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 3F $ $ $ $ $ $ $ $ $ $ $ $ $ $ $ $
000 ***** SYMBOL EXPLANATIONS *****
000 + ONLINE @ PATH NOT VALIDATED - OFFLINE . DOES NOT EXIST
000 * PHYSICALLY ONLINE $ PATH NOT OPERATIONAL

```

Figure 29. Example of DISPLAY M=CHP Command in the MVS/XA and MVS/ESA Environments. Display the status of all devices on channel path 43

Issuing the DISPLAY M=DEV(nnn) command (where nnn is replaced with a specific device number) in the MVS/XA and MVS/ESA environment, you can examine all paths to a device. Figure 30 illustrates information you might see if you issued the command to device 380.

```

000 DISPLAY M=DEV(380)
000 IEE174I 12.45.11 DISPLAY M 753
000 DEVICE 380 STATUS=ONLINE
000 CHP          43 63 44 64
000 PATH ONLINE      Y Y N N
000 CHP PHYSICALLY ONLINE Y Y Y Y
000 PATH OPERATIONAL Y Y N N

```

Figure 30. Example of DISPLAY M=DEV Command in the MVS/XA and MVS/ESA Environments. Display the status of all paths to a specific device.

The DISPLAY M Command in MVS/370

Use the DISPLAY M command to obtain information about the paths to a set of devices. Figure 31 on page 80 is an example of a DISPLAY M command issued in an MVS/370 operating system environment.

You may not be able to determine the exact state of the path by examining the response to a DISPLAY M command without also knowing the state of the device.

For example (see Figure 31 on page 80), if the device is offline, there is no way to tell which paths will come online when the device is varied online.

```

0000 14.20.53      D M=7
0000 14.20.54      IEE0711 14.20.54 DEVICE STATUS 197
197  0 1 2 3 4 5 6 7 8 9 A B C D E F
197 70 *****
197 71 *****
197 72 *****
197 73 *****
197 74 *****
197 75 *****
197 76 *****
197 7E *****
197 7F *****
0000 14.21.14      W A

```

Figure 31. Example of DISPLAY M Command in the MVS/370 Environment. Display the status of all device paths on channel 7.

The DISPLAY Command with ALLOC Parameter

To determine the jobs currently using a volume, issue a DISPLAY U,,ALLOC,Da,Dn command (where a and n are the device address and number of devices respectively) before issuing a VARY command. A 3380 device cannot be varied offline successfully if there are outstanding allocations against the address being varied offline.

Note: If a service representative intends to perform maintenance or repairs on:

- A 3380 string:
Verify allocation usage for all devices on the string.
- A 3380 controller:
If both controllers are online, one controller can be varied offline without severing access to devices on the string.
- A 3380 head and disk assembly (HDA) or a device:
Verify allocation usage for both devices on the HDA.

Figure 32 shows information returned when the DISPLAY command with U and ALLOC parameters is executed. In this example, the job, copylib, is using devices 281 and 284. This job must end before either device can be taken offline.

```

D u,,alloc,280,6
COPYLIB CREATE IEFBR14 0000
IEE106I 13.22.44 UNITS ALLOCATED 441
UNIT JOBNAME ASID JOBNAME ASID JOBNAME ASID JOBNAME ASID
280 *SYSTEM 0000 VTAM 000F
281 COPYLIB 000E
284 COPYLIB 000E
500 *CONSOLE 0007
060 JES2 000A
068 *SYSTEM 0000 *SYSYEM 0001 JES2 000A TCAS 000B
068 VTAM 000F

```

Figure 32. Example of D u,,alloc Command

Using the DEVSERV Command in MVS/XA and MVS/ESA

If your system includes MVS/XA DFP Version 2 or MVS/DFP Version 3 operating systems, you can investigate hardware or configuration problems with the DEVSERV command. Issuing the DEVSERV command causes the system to issue an I/O request on paths to a device or devices. The resulting console display reflects the current physical state of the path(s). The DISPLAY command returns information from the MVS/XA or MVS/ESA control blocks and may not be as current or accurate as the information returned via the DEVSERV command.

The DEVSERV command returns the following information for one or more devices:

- Logical mode of the device
- Number of data sets allocated on the volume
- Volume serial number
- Channel path ID
- Status of the path

Figure 33 shows the DEVSERV command issued to determine the status of devices 281-287, and 290. The example shows devices 281-285 online; 286, 287, and 290 are offline. Path 27 is available to 281-285, while paths 25, 65, and 67 are not available.

```
DEVSERV PATHS,281,8
IEE459I 09.58.42 DEVSERV PATHS 372
UNIT DTYPE M CNT VOLSER CHPID=PATH STATUS
281,3380J,0,000,338001,25=< 27=+ 65=< 67=<
282,3380J,0,000,338JT1,25=< 27=+ 65=< 67=<
283,3380J,0,000,338JT2,25=< 27=+ 65=< 67=<
284,3380J,0,000,338JT3,25=< 27=+ 65=< 67=<
285,3380K,0,000,338KT1,25=< 27=+ 65=< 67=<
286,3380 ,F,000, ,UCB NOT CONNECTED, PATH STATUS UNAVAILABLE
287,3380 ,F,000, ,UCB NOT CONNECTED, PATH STATUS UNAVAILABLE
290,3380 ,F,000, ,UCB NOT CONNECTED, PATH STATUS UNAVAILABLE
***** SYMBOL DEFINITIONS *****
A = ALLOCATED          F = OFFLINE
O = ONLINE             + = PATH AVAILABLE
< = PHYSICALLY UNAVAILABLE
```

Figure 33. Example of DEVSERV Command in the MVS/XA and MVS/ESA Environments

You can find the complete explanation of the DEVSERV command in appropriate MVS/XA or MVS/ESA *System Commands* manual.

Using the VARY Command

The VARY command is a system operator command, and can be issued through an MVS system control program.

The VARY command is used to make a device or path available (online) or unavailable (offline) to the operating system.⁸

⁸ The exact form of the VARY command differs somewhat, based on the processor involved and the level of operating system used.

Note: If a 3380 device is shared by multiple systems, the VARY command must be issued from each system that you wish to access the device or remove access to the device.

Vary Online/Offline Procedure

You **must** use the following procedure when changing the status of your 3380 devices.

To bring a 3380 online:

1. Set all applicable Enable/Disable switches to the Enable position.
2. Issue the VARY PATH online command from each system that you want to access the 3380.
3. Issue a VARY nnn,ONLINE from each system you want to access the 3380.
4. Issue a MOUNT nnn command from each system you want to access the 3380.

To take a 3380 offline:

1. Issue the VARY nnn,OFFLINE. This tells the system to mark that device for deallocation. MVS will deallocate the device at the next invocation of device allocation. You can initiate this by running a job with a dummy allocate statement. When message "IEF281I nnn OFFLINE ..." is issued, proceed to the next step.
2. Issue the VARY PATH offline from each system accessing the device. Message "IEE303I Path nnn OFFLINE" will be issued when the path goes offline.
3. Set the device Enable/Disable switch to the Disable position.

Note: If you do not vary the device offline properly before setting the Enable/Disable switch on the 3380 to the Disable position, a false equipment check error message may be generated. (The system message is a permanent I/O error message indicating an equipment check error condition.) If you want to take a 3880 or 3990 offline, all paths (from all systems) through the 3880 storage directors or 3990 storage paths to all devices accessible by the 3880 or 3990 need to be varied offline. If any of the paths to the devices accessible by the 3880 or 3990 are not varied offline, an **array out-of-sync** condition (see "Out-of-Sync Conditions" on page 84) might result, causing Device End Missing or Start Pending conditions.

The vary process is a **single thread** process: Only one VARY command at a time can be processed by a system. There is no point in entering another VARY command until the current one completes, either successfully or unsuccessfully. Sometimes the vary process takes several minutes, depending on the system activity and what the command has to do. Several forms of the VARY command attempt retries if at first not successful, and this also takes additional time. If the operator enters other VARY commands before the current one completes, the other commands queue behind the current one and may cause more problems later if the current VARY command does not complete successfully. Path-related VARY commands are serialized until the path goes offline. **Exception:** Issue all VARY nnn,OFFLINE (where nnn represents a 3380 device) commands at the same time. Devices that do not immediately go offline do not prevent other VARY commands from processing.

The DISPLAY M command uses the same resources that the VARY command uses. While a VARY command is in process, an operator cannot expect to get a reply from entering a DISPLAY M command.

The VARY Command for Devices

A VARY nnn command applies to a specific 3380 device. (Use VARY nnn-*nnn* to specify a set of devices.) A VARY nnn,ONLINE command allows the system to access the specified 3380 device. A VARY nnn,OFFLINE command prevents the system from using the specified 3380 device from any path of the system that issued the VARY command.

Note: Use the procedures explained in “Vary Online/Offline Procedure” on page 82 to change the status of devices.

The syntax of the VARY command used to vary a device online or offline is:

VARY nnn,ONLINE

VARY nnn,OFFLINE

nnn (three hexadecimal digits) represents a specific device.⁹

The VARY nnn,OFFLINE command will not take a device offline until all allocations against the device (from the system processing the VARY command) have been removed. A device that contains system data cannot be put in offline status with a VARY OFFLINE command, because the device is permanently allocated to the operating system. For devices that do not contain system data, it may be necessary to cancel jobs to put the device in an offline condition.

In a multiple processor environment, activity against a 3380, 3880, or 3990, can affect all the processors. Before a unit is worked on, the system operator should vary offline, from **all** processors, any attached device or control unit. Failure to do this could cause the Device End Missing and Start Pending conditions.

The VARY PATH Command

A VARY PATH command applies to all components of a path, including a channel, a storage director in a 3880 or a storage path in a 3990, and a 3380 controller to the specified 3380 devices. A VARY PATH ONLINE command allows the system to attempt to use the specified 3380 string controller. A VARY PATH OFFLINE command prevents the system from attempting to use the specified 3380 string controller as a path to the specified device.

The syntax of the VARY PATH commands used in the MVS/ESA, MVS/XA and MVS/370 environments follow:

VARY PATH(nnn,pp),ONLINE

VARY PATH(nnn,pp),OFFLINE

nnn identifies a 3380 device accessible through the storage control.

pp is an MVS/ESA or MVS/XA channel path ID.

VARY PATH(nnn,mm),ONLINE

VARY PATH(nnn,mm),OFFLINE

nnn identifies a 3380 device accessible through the storage control.

mm identifies a System/370 channel set.

Note: When a path is varied offline, it can only be brought back online with the VARY PATH ONLINE command.

⁹ In the MVS/ESA or MVS/XA environment, substitute the device number for nnn. In the MVS/370 environment, substitute the unit address for nnn.

If one string controller of a 3380 Model AA4, AD4, AE4, AJ4, or AK4 is not usable because of an equipment problem, the other controller can be selected to access a device on the same string. Vary the unusable controller offline to avoid continuing attempts by the system to use it.

The VARY PATH Command and Path Group ID

The VARY command maintains path group IDs in the arrays used for dynamic path selection functions. When you remove a path from the array with the VARY command, corresponding information is available to the operating system.

Note: To ensure that the 3380 and storage control internal arrays stay in sync with the operating system (and to prevent a data integrity exposure), whenever any physical action is going to be taken on either a 3380 or a storage control, use the procedure in "Vary Online/Offline Procedure" on page 82 to vary a device offline (or vary all devices offline if a storage control is being disabled). Use the DISPLAY U command to verify that a unit is offline before setting an Enable/Disable switch to the Disable position.

Out-of-Sync Conditions

When an internal array does not match what the unit control block (**UCB**) in the operating system shows the array state to be, an **out-of-sync condition** exists. An out-of-sync condition can be caused by an equipment failure in the 3380 and/or storage control but is usually caused by an operator action (such as setting an Enable/Disable switch to Disable) that causes a reset to the storage control's channel interface.

If the UCB in the operating system is not changed before resetting the channel interface, the operating system might repeatedly attempt to access a reserved device over a path that no longer has an eligible path group ID in the array. In this case, the device will appear busy even when it is not in use.

In an MVS/ESA or MVS/XA operating system environment, the subsystem also will not be able to implement dynamic reconnection on all the possible paths that would otherwise be available. This means that the array in the 3380 and/or storage control that contains the path group IDs is out-of-sync with information used by the operating system (the UCB).

Correcting an Out-of-Sync Condition

MVS operating systems have a missing interrupt handler. If no status is returned to the operating system by a device after a set time, the missing interrupt handler redrives the I/O by generating a dummy unit check. In MVS/ESA or MVS/XA, the missing interrupt handler rebuilds the internal arrays.

If an out-of-sync condition exists, the system operator can take several actions to identify the problem and to restore synchronization. For system-related reserve, messages help determine that a problem exists. When a device is not reserved, there is no explicit indication that the array is out-of-sync.

The first indication of a possible out-of-sync problem will probably be a message:

- For MVS/370 operating systems, a Device End Missing system console message
- For MVS/ESA or MVS/XA operating systems, a Start Pending system console message.

If messages issued as a result of the DISPLAY UNITS system operator command are A-BSY-R or O-BSY-R, suspect an out-of-sync problem. If an out-of-sync problem exists, the Device End Missing or Start Pending message will persist. Subsequent DISPLAY UNITS commands will continue to show A-BSY-R or O-BSY-R.

If an out-of-sync condition exists that is not corrected by the system:

1. Issue a VARY PATH, ONLINE for each path to the device
2. Issue a VARY nnn, ONLINE for the device

This will refresh the internal array for all online paths to all devices from the originating system, and thus restore synchronization. You can assume successful restoration of synchronization if the Device End Missing or Start Pending messages stop appearing.

Alternatively, an IPL of the operating system re-creates the internal array and restores synchronization between the array and the IPLed operating system UCB.

If the array is out-of-sync because of a hardware problem, a service representative should be notified to repair the malfunctioning unit. After repair, the array can be re-created by using the VARY (nnn),ONLINE command or by performing an IPL of the operating system.

Missing Interrupts

There is no attention button on the 3380. Missing interrupts cannot be resolved by generating an attention interrupt. With MVS operating systems, missing interrupt conditions are handled by the operating system without operator intervention.

Chapter 8. Moving Data in the MVS Environment

Before you begin moving data from one volume to another, make backup copies of your data. This chapter describes how to backup a volume (to tape), restore a full volume, and restore selected data sets.

When moving DASD data between unlike devices, you can use the DFDSS commands: data set COPY or DUMP/RESTORE. When moving DASD data between like devices, you can use the DFDSS commands: data set COPY, full-volume COPY, or DUMP/RESTORE.

There are many tools, as described in “Planning for Software Tool Availability” on page 11, to move data in the MVS environment. This does not mean that you cannot use other programs to achieve equivalent results, or that you can use DFDSS only in the ways shown.

For information about invoking DFDSS commands from ISMF panels, see the appropriate MVS ISMF User’s Guide. For more detailed information on moving data with DFDSS, see *DFDSS: User’s Guide and Reference*.

Appendix A, “Sample DASD Installation and Migration Project Plan” on page 121 contains a checklist, Moving Data, that you can use to track the completion of this activity.

Modifying Examples for Your Use

The examples in this chapter use the filtering capability of DFDSS. For example, by specifying `DATASET(INCLUDE(USER1.**))`, DFDSS will move all supported data sets with a high-level qualifier of USER1. Use qualified names that are appropriate to your needs.

A JOBCAT or STEPCAT DD statement for the user catalog is required in these examples if either:

- A full-volume COPY operation (such as shown in Figure 38 on page 96) is being performed, and the volume contains a non-SMS managed VSAM data set that is cataloged in a VSAM user catalog or
- A data set COPY operation is being performed, and some of the non-SMS managed data sets are cataloged outside the standard catalog search order.

All the examples in this chapter assume that the user has the appropriate RACF or password authorization. DFDSS checks authorization on DASD volumes and data sets to ensure that you have sufficient access authority. The required access authority depends on the function being performed. If the source¹⁰ data sets are RACF-protected and if you are authorized to access them, DFDSS will define the target data sets to RACF after the move. For details about how DFDSS checks authorization and defines the RACF profile for the target data sets, see *DFDSS: User’s Guide and Reference*.

¹⁰ Source refers to the data set or volume to be copied; target refers to where the copy is placed.

Backup and Recovery While Moving Data

After you have installed and prepared the new units, you will be able to move all or part of your installation's existing data to them. Before moving any data, make backup copies of it, in case it is necessary to restore the data to its original state. For additional information and examples describing logical and physical data set dump and restore processing, see *DFDSS: User's Guide and Reference*.

Creating a Backup Tape

Obtain an IEHLIST VTOC listing and an access method services IDCAMS LISTCAT listing of the data sets before and after moving the data. These listings will be useful for tracking and recovery purposes.

If you are moving groups of data sets (for example, by data set, or by partially qualified data set name), or entire volumes, you may want to back up the appropriate data sets instead of the entire volumes. If you do not specify input volumes, DFDSS will locate input data sets via the standard catalog search order. In this case, a *logical* dump is performed.

The COMPRESS parameter of the DUMP command uses tape more efficiently and reduces tape mount activity. However, you should note that the stand-alone restore program cannot be used to restore from tapes produced with the COMPRESS parameter. Therefore, if you plan to recover your system residence volume with the stand-alone restore program, do not DUMP your SYSRES volume with the COMPRESS parameter.

Restoring a Volume

If it becomes necessary to restore a volume to its original state, you can do so using your backup tape. The PURGE parameter of the RESTORE command allows DFDSS to restore the volume from tape, even if the DASD volume contains unexpired data sets.

Restoring Selected Data Sets

If it is necessary to restore only a data set or a group of data sets, you can use your backup copy, created by either a volume DUMP or a data set DUMP. The REPLACE parameter allows DFDSS to overlay an existing data set of the same name with a data set from the backup tape.

Note: If you restore data sets using a logical data set dump and you are using DFDSS Version 2 Release 2 or higher, you can restore to a volume on an unlike device, capacity permitting. You can also restore data sets from a physical data set dump, but only to a volume on a like device.

Using the DFDSS COPY Command

Data can be moved or copied with DFDSS; moved data is deleted from the source volume. Moving data requires the DELETE parameter on the COPY command. In addition, unexpired data sets are moved only if you specify both the DELETE and PURGE parameters. If you do not specify the DELETE parameter on the COPY command, the original data is left intact and copied to the target volume. However, if the original data is cataloged, the target data must either be cataloged in a different catalog using RECATALOG(newcatname) or renamed using RENAMEUNCONDITIONAL.

At job completion, DFDSS prints a list of successfully processed data sets, a list of data sets that failed to get processed, and a list of data sets that encountered postprocessing errors such as cataloging or redefining data sets to RACF. You then need to refer to previous messages to determine the specific errors and the necessary corrective action. For a description of these messages, see *DFDSS: User's Guide and Reference*.

Data Set COPY

To copy individual data sets or groups of data sets, you specify the DATASET parameter on the COPY command. In addition to selecting data sets based on fully or partially qualified names, you can also select data sets based on data set allocation characteristics. This selection process known as filtering, is described in the *DFDSS: User's Guide and Reference*.

Cataloged Data Set Support

The DFDSS full-volume COPY command can copy VSAM data sets that are cataloged in integrated catalog facility or VSAM catalogs. The DFDSS data set COPY command supports VSAM data sets that are cataloged in integrated catalog facility catalogs, and also supports non-VSAM data sets. However, catalog entries will not be updated.

Catalog Search Order

If multiple copies of data sets exist on your system, it is important to understand the standard catalog search order. This order determines which copy is found and used by DFDSS. For a complete description of the standard catalog search order, see the *MVS/ESA or MVS/XA Integrated Catalog Administration: Access Method Service Reference*.

Moving Catalogs with DFDSS

You should move catalogs in an independent step and make backup copies of any catalogs before moving them, in case a hardware or user error occurs. Move user catalogs only in a controlled environment. Do not move a catalog together with the data sets contained in it. Master catalogs cannot be moved with DFDSS; use IDCAMS function to move these catalogs.

You can use the MVS/XA DFP or MVS/DFP catalog locking facility (described in Chapter 5, "Planning the Data Configuration" on page 47) to lock catalogs when you move them. When you lock a catalog, DFDSS recognizes the locked status and maintains it while moving the catalog. After the catalog has been moved, you can unlock it for normal use.

When moving an integrated catalog facility user catalog, you must always specify its fully qualified data set name. Since the catalog cannot be renamed or

uncataloged on the target volume, the DELETE parameter is required to delete the catalog from the source volume.

When moving CVOL catalogs, you must specify the DELETE parameter. The RENAMEUNCONDITIONAL and UNCATALOG parameters are not allowed, and the CATALOG parameter is the default. DFDSS invokes access method services or OS utilities to move catalogs. You need to rebuild catalog aliases after moving a catalog, unless you are using MVS/XA DFP 2.3.0 or later, or MVS/DFP Version 3.1.0 or later, with DFDSS Version 2.3.0, or later, which rebuilds the catalog aliases for you.

In Figure 34, user catalog USERCAT1 is moved from a 3350 volume to a 3380 volume. The catalog name, USERCAT1, must be explicitly specified in the INCLUDE dataset list. The RECATALOG(*) parameter records the new location of USERCAT1 in the same catalog that pointed to the old location of USERCAT1. Do not move catalogs unless you specify the UTIL MSG=YES parameter as illustrated in the example. The DELETE parameter removed the reference in the catalog to the old location.

```
/*
//DFDSS EXEC PGM=ADRDSSU,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//DD1 DD UNIT=3380,VOL=SER=3380L1,DISP=OLD
//SYSIN DD *
COPY OUTDD(DD1) DATASET(INCLUDE(USERCAT1)) RECATALOG(*) DELETE
/*
```

Figure 34. Using DFDSS to Move an Integrated Catalog Facility User Catalog

Access method services can also be used to move catalogs. The access method services REPRO command copies catalogs, splits integrated catalog facility catalog entries between two catalogs, and merges integrated catalog facility entries into another integrated catalog facility user catalog or master catalog. For more information, see the access method services reference manual for your system.

Moving Unmovable Data Sets with DFDSS

When copying unmovable data sets to like devices, DFDSS will place them at the same track locations on the target volume if all the following conditions are true,

- The operating system is MVS/XA DFP Version 2 Release 1.0, MVS/DFP Version 3 Release 1.0, or any later release.
- If the data set is an ISAM data set, the target volume has an indexed VTOC.
- The space where the unmovable data set would reside is available.

If any of these conditions is false, DFDSS will not copy an unmovable data set unless the FORCE parameter is specified. The FORCE parameter enables DFDSS COPY to treat unmovable data sets as movable data sets and also to move them between unlike devices. When the FORCE parameter is used, DFDSS will place the specified data sets in any available location on the target volume. Consequently, use FORCE with caution.

If some data sets have CCHHR (cylinder, cylinder, head, head, record) location-dependent data, and you have specified FORCE to move other, location-independent unmovable data sets, use the EXCLUDE parameter to prevent movement of the location-dependent data sets.

Another way to position data sets in a specific location on a volume is to allocate all space on the target volume, except where the unmovable data sets will be placed. Then move the unmovable data sets with the FORCE parameter, after which you can scratch the dummy space allocation.

In Figure 35, USER1 moves unmovable data sets that begin with USER1.UNMOVA to volume 338NE1 from 3350 volumes 335OL1 and 335OL2 with any overflow data sets moved to volume 338NE2. When 338NE1 is full, the PERCENTUTILIZED parameter directs DFDSS to use as much as 20% of the capacity of 338NE2. Two of the data sets, USER1.UNMOVA.DATASET8 and USER1.UNMOVA.DATASET9, are excluded from the move operation. All source data sets, including unexpired data sets, that are selected are deleted. The target data sets are cataloged in the same catalog that points to the source data sets.

```

//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//DD3 DD VOL=SER=335OL1,UNIT=3350,DISP=OLD
//DD4 DD VOL=SER=335OL2,UNIT=3350,DISP=OLD
//DD1 DD VOL=SER=338NE1,UNIT=3380,DISP=OLD
//DD2 DD VOL=SER=338NE2,UNIT=3380,DISP=OLD
//SYSIN DD *
COPY INDD(DD3,DD4) OUTDD(DD1,DD2) -
DATASET(INCLUDE(USER1.UNMOVA.**)) -
EXCLUDE(USER1.UNMOVA.DATASET8,USER1.UNMOVA.DATASET9)) -
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP PURGE -
FORCE PERCENTUTILIZED(100,20)
/*

```

Figure 35. Using DFDSS to Move Selected Unmovable Data Sets

Moving Multivolume Data Sets

You can copy cataloged multivolume data sets to a single volume or to multiple volumes. Volume spread is handled by DFDSS as follows:

- For non-VSAM data sets, DFDSS merges the data set into a single volume, if that volume has sufficient space.
- For VSAM key sequence data sets that have the index and data components on different volumes, DFDSS preserves the volume spread if enough target volumes are specified.
- For VSAM data sets where a component spans multiple volumes, the multivolume component is merged into a single volume, if that volume has sufficient space.
- For VSAM key range data sets, where there are as many volumes as key ranges, DFDSS preserves the volume spread if enough target volumes with sufficient space are specified.

If you move a multivolume data set containing standard user labels, then only the standard user label on the first volume is copied to the target volume.

To copy any part of a multivolume data set, you must copy the entire data set. DFDSS Version 2.3.0 provides the ALLMULTI (ALLM), LOGINDDNAME (LIDD), and LOGINDYNAM (LIDY) parameters to use when copying or dumping multivolume data sets. The ALLM parameter identifies the data set as a multivolume data set. At least one of the volumes on which it resides must be specified with the LIDD or LIDY parameter. To move a multivolume data set with earlier releases of DFDSS,

you specify its name and **all** the volumes on which it resides, using INDDNAME (INDD) or INDYNAME (INDY)¹¹.

Figure 36 shows you how to move a multivolume data set named USER1.DATA, from two volumes (335001 and 335002) to one volume (3380L1). Another output volume (3380L2) is specified for possible overflow. Because INDD is used, **all** input volumes where the data set resides must be specified on separate DD statements. When DFDSS invokes IEHMOVE to copy multivolume non-VSAM data sets, all input volumes must be specified in the first DD statement. The example accommodates both requirements.

```
/*
//DFDSS EXEC PGM=ADROSSU,PARM='UTILMSG=YES'
//STEP CAT DD DSN=USERCAT1,DISP=SHR
//SYSPRINT DD SYSOUT=*
//DD1 DD VOL=SER=(335001,335002),
// UNIT=(3350,2),DISP=OLD
//DD2 DD DISP=OLD,VOL=SER=335002,UNIT=3350
//TODD1 DD DISP=OLD,VOL=SER=3380L1,UNIT=3380
//TODD2 DD DISP=OLD,VOL=SER=3380L2,UNIT=3380
//SYSIN DD *
COPY INDD(DD1,DD2) -
OUTDD(TODD1,TODD2) -
DATASET(INCLUDE(USER1.DATA)) -
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP -
PERCENTUTILIZED(100,20) PURGE
/*
```

Figure 36. Using DFDSS to Move a Multivolume Data Set

Moving VSAM Data Sets

The DFDSS data set copy operation supports VSAM moves done at the cluster level. KSDS (key sequence data sets) including those within the ranges specified, ESDS (entry sequence data sets), RRDS (relative record data sets), and LDS (linear data sets) are supported.

When moving a VSAM data set with alternate indexes, you must move the alternate indexes independently of the base cluster.

When moving a VSAM alternate index cluster, RENAMEUNCONDITIONAL and RECATALOG(newcatname) are not allowed and DELETE is required; the moved alternate index cluster will continue to be related to its base cluster.

Note: When moving a VSAM data set, the CATALOG parameter is used as the default and UNCATALOG is ignored. The catalog must be in the standard catalog search order.

For password-protected VSAM data sets in password-protected catalogs, you can specify the password of the catalog instead of the password of each data set. If you specify both the data set and the catalog passwords, only the data set password is used to check authorization.

¹¹ Similarly, OUTDD is an abbreviated form of OUTDDNAME, and OUTDY is an abbreviated form of OUTDYNAM.

Moving IMS, CICS and DB2 Data Bases

To maintain the integrity of these data bases, use the appropriate data base utility provided to move these data base data sets. For more information, see Chapter 9, "Moving IMS/VS and DB2 Data Sets" on page 103.

Moving SAM/PAM Data Sets and Data Sets with Null DSORG

Unless you specify the ALLDATA and ALLEXCP parameters, only the used (rather than allocated) tracks¹² are copied for SAM data sets and PAM data sets; data sets with unknown data set organization are not copied.

Moving ISAM Data Sets

When you move an ISAM data set, the target data set has the prime, index, and overflow areas combined into one allocated area. You can also use the IEBISAM utility COPY command to move ISAM data sets.

Moving BDAM Data Sets

DFDSS Version 2.3.0 and later releases move all BDAM data sets to like and unlike devices. Unless you specify the RELBLOCKADDRESS(dsn [...]) parameter, DFDSS assumes BDAM data sets are organized such that their records are located using TTR.¹³ DFDSS invokes IEHMOVE when moving TTR-organized data sets to an unlike device; keep in mind that the target volume must have an equal or larger capacity than the source volume.

If a BDAM data set in TTR format has location dependent data (for example, if a CCHHR addressing scheme is used) and the data set does not have the unmovable attribute, you can prevent DFDSS from moving it by specifying its name in the EXCLUDE parameter.

When you specify the RELBLOCKADDRESS(dsn [...]) parameter, DFDSS assumes that the data sets specified by the RELBLOCKADDRESS parameter are organized by relative block address (and not by TTR), and moves them without maintaining the relative track and record number for each record. They will be moved block by block, provided that the blocks fit on the tracks of the target device.

Unmovable BDAM data sets will not be moved unless you specify the FORCE parameter. For information about moving unmovable data sets, see "Moving Unmovable Data Sets with DFDSS" on page 90.

Moving System Data Sets

Not all system data sets need moving; some are allocated during the system generation process, others are built at IPL time. However, other system data sets can be moved by DFDSS. For a listing of these data sets see Chapter 5, "Planning the Data Configuration" on page 47.

Unless you exclude them, system data sets will be copied. They generally remain open while the system is running, and thus cannot be scratched or uncataloged. DELETE and UNCATALOG operations are supported only for those data sets that are not in use. Any data sets that have a high level qualifier of SYS1 will not be

¹² "Used tracks" consist of the tracks from the beginning of a (SAM or PAM) data set to the last-used track, as indicated in the DS1LSTAR field of the format-1 DSCB.

¹³ TT is the relative track number from the beginning of the data set, and R is the record number on the track.

scratched and uncataloged. System data sets can be moved by DFDSS, using one of the following methods:

- Dump the data sets and restore them to a new volume.
- Copy them to another volume and catalog them in a different catalog, as shown in Figure 37.

In Figure 37, the TOLERATE(ENQFAILURE) parameter specifies that data sets are to be processed even though shared or exclusive access cannot be obtained. If you specify this optional parameter, DFDSS is more likely to gain access to them, but it is possible that they may be updated by other applications during the copy process. (Note that, on data set copy operations in which a utility is invoked by DFDSS to move the data, DFDSS ignores the TOLERATE(ENQFAILURE) parameter. Thus, data sets that are in use will not be copied.) The WAIT(1,1) parameter specifies the number of seconds DFDSS is to wait between attempts to obtain control of a data set, and the number of times to retry the operation. Use the DUMP/RESTORE method of moving system data sets to avoid these problems.

```
//*
//DFDSS EXEC PGM=ADRDSSU
//STEP CAT DD DSN=USERCAT1,DISP=SHR
//SYSPRINT DD SYSOUT=*
//DD3 DD DISP=OLD,VOL=SER=MXARES,UNIT=SYSDA
//DD4 DD DISP=OLD,VOL=SER=SPOOL1,UNIT=SYSDA
//DD1 DD DISP=OLD,VOL=SER=3380L1,UNIT=SYSDA
//DD2 DD DISP=OLD,VOL=SER=3380L2,UNIT=SYSDA
//SYSIN DD *
COPY INDD(DD3,DD4) OUTDD(DD1,DD2) RECATALOG(USERCAT1)
DATASET(INCLUDE(SYS1.**),EXCLUDE(SYS1.VVDS.**),SYS1.VTOCIX.**))
PERCENTUTILIZED(90,25) ALLDATA(*) ALLEXCP
TOLERATE(ENQFAILURE) WAIT(1,1)
/*
```

Figure 37. Using DFDSS to Move System Data Sets

When a DFDSS data set copy operation is used to copy the following data sets, space is allocated for the target data set, but no data is copied:

- Model DSCBs
- Page and swap data sets
- SYS1.STGINDEX

For additional information about the DFDSS COPY command, see *DFDSS: User's Guide and Reference*.

Moving SMS-Managed Data Sets

In an SMS environment, the DFDSS COPY command invokes ACS routines which determine candidate target volumes for the output data sets.

In most cases, you allow the system to control placement of data sets. However, SMS allows you to override the storage classes assigned by the ACS routines and place a data set on a specific volume.

If you want to control output data set allocation to specific volumes, you can use ACS routines and storage class. If a data set's storage class has the guaranteed space attribute, the data set is placed on the volumes specified by the DFDSS OUTDDNAME or OUTDYNAM parameters, if the volumes reside in the same storage group. By specifying BYPASSACS and STORCLAS parameters, you can

ensure that a storage class which has the guaranteed space attribute will be reassigned. Thus, the volumes you specify with OUTDDNAME or OUTDYNAM are used as target volumes.

For additional information about using the DFDSS COPY command in an SMS environment, see *DFDSS: User's Guide and Reference*.

Full-Volume COPY

If you do not specify the DATASET or TRACKS parameter on the COPY command, or if you specify FULL, the full volume will be copied. In this case, you must specify the INDD (or INDY) parameter to indicate what volume is to be copied, and the OUTDD (or OUTDY) parameter to indicate the target volume. This operation copies data from the source volume to the target volume in track format. Consequently, it can only be performed on like devices of equal or greater capacity; for example, from a standard capacity 3380 to any 3380 model, from a double capacity 3380 model to a double or triple capacity 3380 model, or from a triple capacity 3380 model to another triple capacity 3380 model.

If you specify the COPYVOLID parameter, the VOLID (the volume serial number) from the source DASD volume is copied to the target DASD volume. Note that, if you change the VOLID on a DASD volume, the VOLID in the system UCB of the target volume is cleared, the operator is notified, and the operating system initiates a demount of the target volume at the end of the COPY operation. To use the target volume, you must demount the source volume and then mount the target volume online.

Because of the special processing requirements of some data sets, if you intend to use the data set COPY operation to move volumes, you must move the data sets in an appropriate sequence to ensure that the expected results are achieved. Examples of this kind of data set are:

- Unmovable data sets
- Multivolume data sets
- Integrated catalog facility user catalogs
- Data sets beginning with SYS1
- Data sets that are used by device-dependent application programs.

You may want to process unmovable data sets first to place them at the same track location on the target device. Move user catalogs only in a controlled environment. **Do not move catalogs together with the data sets cataloged in them.** Also note that some data sets are ineligible for copy by DFDSS (for example, nonintegrated catalog facility VSAM data sets), and some data sets require special parameters (for example, unmovable data sets).

VTOC Considerations

Make certain the VTOC on the target volume is large enough to contain all the entries for the data sets to be placed on it. A DFDSS full-volume copy copies the VTOC from the first source volume specified to the target volume. If you determine that this VTOC will not be large enough to contain the entries for the data sets combined on the target volume from the source volumes, use ICKDSF to reinitialize the target volume with a larger VTOC (see "Calculate VTOC Size" on page 74). Then use DFDSS data set COPY to transfer the data sets.

With DFDSS Version 2.3.0, you can take advantage of the performance benefits gained from using DFDSS full volume COPY when moving data to a large capacity volume. If you allocate a VTOC on a larger capacity volume, DFDSS Version 2.3.0

(with the appropriate PTF), allows you to provide an expanded VTOC while using DFDSS full-volume COPY to migrate the data. In this case, the full volume COPY function will recognize and use an expanded VTOC and VTOC index allocated on a larger capacity target volume. First, use ICKDSF to initialize the larger capacity volume. Place the expanded VTOC and expanded VTOC index at a track address beyond the track address range of the lesser capacity source volume. Then use DFDSS full volume COPY to move the data to the larger capacity volume. Full volume COPY will update the expanded VTOC and expanded VTOC index to reflect the data sets, and additional freespace on the larger capacity volume.

One 3380 Volume to Another 3380 Volume

An easy way to move data to a 3380 volume is with the DFDSS full-volume COPY command as shown in Figure 38. The ALLDATA(*) and ALLEXCP parameters cause all allocated space on the source volume to be moved.

DFDSS parameter COPYVOLID causes the volume's VOLID to be copied to the target volume. (Note that, after the full-volume COPY operation, the operating system will initiate a demount of the target volume. To use the target volume, you must demount the source volume and then remount the target volume.)

```
/*  
//DFDSS EXEC PGM=SPIDSSU  
//SYSPRINT DD SYSOUT=*  
//SYSUDUMP DD SYSOUT=V  
//DD1 DD DISP=OLD,VOL=SER=338002,UNIT=SYSDA  
//DD2 DD DISP=OLD,VOL=SER=3380L1,UNIT=SYSDA  
//SYSIN DD *  
COPY INDD(DD1) OUTDD(DD2) ALLDATA(*) ALLEXCP COPYVOLID  
/*
```

Figure 38. Using DFDSS to Copy by Volume from 3380 to 3380. The target volume must be the same size or larger than the source volume.

The target volume in Figure 38 must be of equal or greater capacity than the source volume. You cannot use this example to move data from a larger capacity 3380 volume to a smaller capacity 3380 volume.

Make certain the VTOC moved with the first volume is large enough to describe all the data sets you intend to place on the target volume.

This operation is performed on a track-for-track basis, allowing you to move all data set types from the source volume.

You can also use this example to copy the data from a standard capacity 3380 volume to a double capacity 3380 volume or triple capacity 3380 volume. Likewise, you can use this example to copy the data from a double capacity 3380 volume to a triple capacity 3380 volume. Source data, including the VTOC, will be copied to the first portion of the larger capacity 3380 volume.

If you want to combine standard capacity 3380 volumes or if you want to move other data sets to the remaining cylinders of the larger capacity 3380 models, you must use the DFDSS data set COPY command.

Three Standard Capacity 3380 Volumes to One Triple Capacity 3380 Volume

There are many ways to move the data from standard capacity 3380 volumes to larger capacity 3380 volumes. One method is to move the first volume with a DFDSS full-volume COPY command, as described in "One 3380 Volume to Another 3380 Volume" on page 96, and then move the other volumes with a DFDSS data set COPY commands. First, if you determine that the VTOC from the first source volume copied is not large enough to hold all the entries for the data sets you wish to place on the target volume, reinitialize the target volume with ICKDSF to create a larger VTOC (see "Calculate VTOC Size" on page 74). Then use DFDSS full volume COPY to move the data from the source volume to the target volume. Move the data on the remaining two 3380 volumes to the target volume using DFDSS data set copy. Another method of moving data from standard capacity 3380 volumes to triple capacity 3380 volumes is to use DFDSS data set copy. To minimize the impact of an abnormal job termination, you may choose to divide the data move into several smaller jobs.

After handling any data sets requiring special processing as described in "Full-Volume COPY" on page 95, you could use the example shown in Figure 39 to move all the data sets belonging to USER1 on the three standard capacity 3380 volumes to one triple capacity 3380 volume with a single DFDSS data set COPY command. If you do not want DFDSS to process specific data sets, be sure to identify them using the EXCLUDE parameter or other filtering criteria. If you want to copy all the data sets on the source volumes, use DATASET(INCLUDE(**)).

Source volumes 338001, 338002, and 338003 are moved to target volume 3380L1. An overflow volume, 3380L2, is provided in case the target volume becomes full during the operation. The PERCENTUTILIZED(100,20) parameter specifies that when 3380L1 is full, DFDSS is to use as much as 20% of the capacity of volume 3380L2. For VSAM data sets targeted for multiple volumes, DFDSS may ignore the PERCENTUTILIZED parameter.

```
/*
//DFDSS EXEC PGM=ADRDSSU,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//DD3 DD DISP=OLD,VOL=SER=338001,UNIT=3380
//DD4 DD DISP=OLD,VOL=SER=338002,UNIT=3380
//DD5 DD DISP=OLD,VOL=SER=338003,UNIT=3380
//DD1 DD DISP=OLD,VOL=SER=3380L1,UNIT=3380
//DD2 DD DISP=OLD,VOL=SER=3380L2,UNIT=3380
//SYSIN DD *
COPY INDD(DD3,DD4,DD5) OUTDD(DD1,DD2)
DATASET(INCLUDE(USER1.**))
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP PURGE
PERCENTUTILIZED(100,20)
/*
```

Figure 39. Using DFDSS to Move Three Standard Capacity 3380 Volumes to a Triple Capacity 3380 Volume

The source data sets, including unexpired data sets, are deleted and the target data sets are recataloged (if they were cataloged in the source volumes). Moved data sets must conform to the standard catalog search order (see "Catalog Search Order" on page 89) to be recataloged.

Three Double Capacity 3380 Volumes to Two Triple Capacity 3380 Volumes

The fastest method for moving data from double capacity 3380 volumes to triple capacity 3380 volumes is by using DFDSS full volume COPY. Full volume copy can move a double capacity 3380 volume to the upper two thirds of a triple capacity volume.

When you move three double capacity 3380 volumes to two triple capacity volumes use the following procedure:

1. Use DFDSS full volume COPY to move two double capacity volumes to the two target triple capacity volumes (upper two thirds each). Ensure that the source data set's catalog and RACF profiles reflect the correct volume serial number by specifying the COPYVOLID keyword.
2. Use DFDSS data set COPY to move data from the third double capacity volume to the lower third of both target triple capacity volumes.

The volume with the least number of data sets should be moved using DFDSS data set COPY because performance is related to the number of data sets moved. Volumes with unmovable or device-dependant data set requirements should be moved using DFDSS full volume COPY.

Note: If the VTOC from the source volume is not large enough to hold all the entries for the data sets you wish to place on a target, before performing a DFDSS full volume copy, you should reinitialize the target volume with ICKDSF to create a larger VTOC.

The following example shows three double capacity source volumes (338001, 338002, 338003) moving to two target triple capacity volumes (338004, 338005). In the first step, volumes 338001 and 338002 are moved using DFDSS full volume COPY. DFDSS varies the target volumes offline since their volids (now 338001, 338002) are not allowed to be online because they are duplicates. You must vary the source volumes offline and vary the target volumes 338004 and 338005 (now labeled 338001 and 338002) back online.

Once the first step is completed, volume 33803 moved to 338001 and 338002 using DFDSS data set COPY. Volume 338002 is provided as an overflow volume and is used only when volume 338001 has reached the desired capacity. The PERCENTUTILIZED(80,80) parameter specifies that when 338001 is filled to 80% capacity, DFDSS is to begin placing data sets on volume 338002.

```

/* JOB 1, Full Volume COPY
//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//DDE1 DD DISP=OLD,VOL=SER=338001,UNIT=3380
//DDE2 DD DISP=OLD,VOL=SER=338002,UNIT=3380
//DDK1 DD DISP=OLD,VOL=SER=338004,UNIT=3380
//DDK2 DD DISP=OLD,VOL=SER=338005,UNIT=3380
//SYSIN DD *

COPY FULL INDD(DDE1) OUTDD(DDK1) -
COPYVOLID
COPY FULL INDD(DDE2) OUTDD(DDK2) -
COPYVOLID

/* JOB 2, Data Set COPY
//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//DDE3 DD DISP=OLD,VOL=SER=338003,UNIT=3380
//DDK1 DD DISP=OLD,VOL=SER=338001,UNIT=3380
//DDK2 DD DISP=OLD,VOL=SER=338002,UNIT=3380
//SYSIN DD *

COPY INDD(DDE3) OUTDD(DDK1,DDK2) -
DATASET(INCLUDE(**)) -
RECATALOG(*) DELETE -
PERCENTUTILIZED(80,80)

```

Figure 40. Using DFDSS to Move Three Double Capacity 3380 Volumes to Two Triple Capacity 3380 Volumes

Alternatively, you can use DFDSS data set COPY to move all the data sets from the three double capacity volumes to the two triple capacity volumes. This method is slower, but can be used if the situation requires it.

Two 3350 Volumes to a 3380 Volume

See the example in "Full-Volume COPY" on page 95 for additional details and data sets that may require special processing. Make certain the VTOC on the 3380 volume is large enough to describe all the data sets from all the volumes.

The example shown in Figure 41 on page 100 moves all the data sets belonging to USER1 on the two 3350 volumes to the 3380 volume with a single DFDSS data set COPY command. If you do not want DFDSS to process specific data sets, be sure to identify them using the EXCLUDE parameter or other filtering criteria. If you want to process all the data sets on the source volumes, specify DATASET(INCLUDE(**)).

In Figure 41, data sets from two 3350 volumes (335001 and 335002) are moved to a 3380 volume (3380L1). A second 3380 volume (3380L2) is specified as an overflow volume. The PERCENTUTILIZED(90,25) parameter specifies that when 3380L1 is approximately 90% full, DFDSS is to use as much as 25% of the overflow volume, 3380L2.


```

/*
//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=V
//DD3 DD DISP=OLD,VOL=SER=335001,UNIT=3350
//DD4 DD DISP=OLD,VOL=SER=335002,UNIT=3350
//DD1 DD DISP=OLD,VOL=SER=3380L1,UNIT=3380
//DD2 DD DISP=OLD,VOL=SER=3380L2,UNIT=3380
//SYSIN DD *
COPY INDD(DD3,DD4) OUTDD(DD1,DD2)
DATASET(INCLUDE(USER1.**))
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP
PERCENTUTILIZED(90,25) PURGE
/*

```

Figure 41. Using DFDSS to Move Two 3350 Volumes to a 3380 Volume

The source data sets (including unexpired data sets) are deleted and the target data sets are cataloged in the same catalog in which the source data sets were cataloged. Moved data sets must conform to the standard catalog search order (see "Catalog Search Order" on page 89) to be recataloged.

Four 3350 Volumes to a Double Capacity 3380 Volume

This DFDSS example moving the data from four 3350 volumes to a double capacity 3380 volume is similar to the one moving the data from two 3350 volumes to a standard capacity 3380 volume. See "Two 3350 Volumes to a 3380 Volume" on page 99. Make certain the VTOC on the 3380 volume is large enough to describe all the data sets from all the volumes. See the example in "Full-Volume COPY" on page 95 for data sets that should be processed ahead of or excluded from the data sets processed in this example.

In Figure 42, all USER1 data sets residing on four 3350 volumes (3350L1, 3350L2, 3350L3, and 3350L4) are moved to one double capacity 3380 volume (338NE1). An overflow volume, 338NE2, is provided in case the primary volume becomes full.

```

//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//DD3 DD VOL=SER=3350L1,UNIT=3350,DISP=OLD
//DD4 DD VOL=SER=3350L2,UNIT=3350,DISP=OLD
//DD5 DD VOL=SER=3350L3,UNIT=3350,DISP=OLD
//DD6 DD VOL=SER=3350L4,UNIT=3350,DISP=OLD
//DD1 DD VOL=SER=338NE1,UNIT=3380,DISP=OLD
//DD2 DD VOL=SER=338NE2,UNIT=3380,DISP=OLD
//SYSIN DD *
COPY INDD(DD3,DD4,DD5,DD6) OUTDD(DD1,DD2)
DATASET(INCLUDE(USER1.**))
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP
PERCENTUTILIZED(90,25) PURGE
/*

```

Figure 42. Using DFDSS to Move Four 3350 Volumes to a Double Capacity 3380 Volume

Data sets moved from the 3350 volumes are recataloged and the data sets (including unexpired data sets) are deleted from the 3350 volumes. Moved data sets must conform to the standard catalog search order (see "Catalog Search Order" on page 89) to be recataloged.

Six 3350 Volumes to a Triple Capacity 3380 Volume

This DFDSS example moving the data from six 3350 volumes to a triple capacity 3380 volume is also similar to the one moving data from two 3350 volumes to a standard capacity 3380 volume. Make certain the VTOC on the 3380 volume is large enough to describe all the data sets from all the volumes. See the example in "Full-Volume COPY" on page 95 for data sets that should be processed ahead of or excluded from the data sets processed in this example.

In Figure 43, all USER1 data sets residing on six 3350 volumes (3350L1, 3350L2, 3350L3, 3350L4, 3350L5, and 3350L6) are moved to one triple capacity 3380 volume (338NE1). An overflow volume, 338NE2, is provided in case the primary volume becomes full.

```
//DFDSS EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//DD3 DD VOL=SER=3350L1,UNIT=3350,DISP=OLD
//DD4 DD VOL=SER=3350L2,UNIT=3350,DISP=OLD
//DD5 DD VOL=SER=3350L3,UNIT=3350,DISP=OLD
//DD6 DD VOL=SER=3350L4,UNIT=3350,DISP=OLD
//DD7 DD VOL=SER=3350L5,UNIT=3350,DISP=OLD
//DD8 DD VOL=SER=3350L6,UNIT=3350,DISP=OLD
//DD1 DD VOL=SER=338NE1,UNIT=3380,DISP=OLD
//DD2 DD VOL=SER=338NE2,UNIT=3380,DISP=OLD
//SYSIN DD *
COPY INDD(DD3,DD4,DD5,DD6,DD7,DD8) OUTDD(DD1,DD2)
DATASET(INCLUDE(USER1.**))
RECATALOG(*) DELETE ALLDATA(*) ALLEXCP
PERCENTUTILIZED(90,25) PURGE
/*
```

Figure 43. Using DFDSS to Move Six 3350 Volumes to a Triple Capacity 3380 Volume

Data sets moved from the 3350 volumes are recataloged and the data sets (including unexpired data sets) are deleted from the 3350 volumes. Moved data sets must conform to the standard catalog search order (see "Catalog Search Order" on page 89) to be recataloged.

Using DUMP/RESTORE to Move Data

You may choose to use DFDSS DUMP/RESTORE to move some of your data. As part of the plan to move your data, you should have made backup copies of your volumes or data sets. If you wish, you can restore the backup copies to the new device, instead of using DFDSS COPY.

If you use DFDSS Version 2 Release 2.0 or later, to produce logical dump data sets (which can physically reside on tape or DASD), you can restore them to like or unlike devices.

Note: You should also consider using DFDSS logical dump processing if the backup copy may have to be restored at a remote site or on another system, where unlike devices are installed. For example, logical dump data sets taken from 3380 volumes can be restored to 3350 volumes.

DFDSS Version 2 cannot restore logical data sets dumped with DFDSS Version 1 to unlike devices.

For more details, see *DFDSS: User's Guide and Reference*.

Moving Data Sets to Smaller Capacity 3380 Models

When moving data volumes from larger capacity 3380 volume to smaller capacity 3380 volume, take into account this capacity difference. In some cases, a partially full large capacity 3380 volume can fit entirely on to a smaller capacity 3380 volume. Consider the size of the VTOC on the source volume. You may have to tailor the VTOC if it is too large for the target volume.

With DFDSS data set COPY, you can move data sets on double capacity 3380 volumes and triple capacity 3380 volumes to 3380 models with less capacity. Using DFDSS filtering capability, you can select and move individual data sets to a specific volume. Alternatively, you can move all the data sets on a volume to several target volumes in a single step, spreading them across these volumes with the PERCENTUTILIZED parameter to control the allocated space on each volume.

Moving Tape Data Sets

The increased storage capacity of 3380 direct access storage may make it feasible for you to move some of your smaller and more frequently used tape data sets to DASD. Moving most tape data sets to 3380 volumes can be accomplished using MVS utilities, and requires no unusual JCL. Tape data sets created by DFDSS can be moved using DFDSS, and those created by access method services IMPORT/EXPORT can be moved by IMPORT/EXPORT.

Figure 44 shows how you could move a sequential tape data set to 3380 volume 3380L1 using IEBGENER. This example catalogs the data set, after moving it to the 3380 volume.

```
//*  
//STEP1 EXEC PGM=IEBGENER  
//SYSPRINT DD SYSOUT=*  
//SYSUT1 DD VOL=SER=TAPE01,UNIT=3400-6,DISP=(OLD,KEEP),  
// LABEL=(1,SL),DSNAME=PROG1T  
//SYSUT2 DD VOL=SER=3380L1,UNIT=3380,DISP=(NEW,CATLG),  
// DSNAME=PROG1T,  
// DCB=(RECFM=FB,LRECL=80,BLKSIZE=6160),  
// SPACE=(6160,(20,10),RLSE)  
//SYSIN DD DUMMY  
/*
```

Figure 44. Moving a Tape Data Set to 3380 Direct Access Storage

Chapter 9. Moving IMS/VS and DB2 Data Sets

This chapter provides information on moving data sets to 3380 volumes in the IMS/VS and DB2 environments. Guidelines and considerations are provided to help you develop a data movement strategy for your IMS/VS or DB2 database(s).

Before moving data belonging to any data base, resolve any data inconsistencies using the appropriate utility. Ensure database integrity by either stopping or bringing down the database.

For detailed information, see the documentation for your particular product (IMS/VS or DB2).

Moving IMS/VS Data

Before you begin the data movement process, you should identify the types of IMS data sets in your data processing center and the tools available to move them. IMS data can be categorized into three different groups:

- IMS data sets that are PDS organization
- IMS data sets that are built by IMS
- IMS Data Bases.

Tools you can use to move each group are listed in the following sections.

Moving PDS Data

Use DFDSS or the MVS utility IEBCOPY to move the following partitioned data sets.

- IMS LOAD libraries
- IMS RESLIB
- IMS ACP library
- IMS PGM library
- IMS DBD library
- IMS PSB library
- IMS FORMAT data set
- IMS MATRIX data set.

Moving Data Sets Built by IMS

The following data sets need to be recreated using IMS Utilities or reallocated on a new device for IMS to use.

LOGGING
Write ahead
IMSMON
QBLKS
SHMSG
LGMSG
MODBLKS
MODSTAT
MSDBCP1
MSDBDUMP
MSDBINIT
IMSTFMFTA
IMSSPA
IMSRDS.

Note: The DBRC RECON data sets can be moved using MVS utilities.

Moving IMS Data Bases

The following data bases can be moved using MVS utilities if you do not re-block the data. If the data base is re-blocked, you **must** use IMS utilities to move the data base.

HSAM
SHSAM
HISAM
SHISAM
GSAM
HDAM
HIDAM
DEDB.

Note: Special considerations should be given to IMS data sets that span volumes.

The following step-by-step approach can be used to move a data base from one device to another.

1. Unload your data base using the existing data base description (**DBD**) and the appropriate unload utility.
2. Recalculate the control interval (**CI**) or block size to maximize use of track space on the new device. Information on calculating CI or block size is contained in *IMS/VS Version 2 Data Base Administration Guide*.
3. Code a new data base description.
4. Rebuild the application control block (**ACB**) if you have ACBs pre-built rather than built dynamically.
5. For non-VSAM data sets, delete the old data base space and define space for the new data base. For VSAM data sets, delete the space allocated for the old cluster(s) and define space for the new cluster(s).
6. Reload your data base, using the new DBD and the appropriate reload utility. Remember to make an image copy of your data base as soon as it's reloaded.
7. If your data base uses logical relationships or secondary indexes, you'll have to run some of the reorganization utilities before and after reloading to resolve

prefix information. A flowchart on which utilities to use and the order in which they must be run is contained in *IMS/VS Version 2 Data Base Administration Guide*.

Using Image Copy/Recovery Utilities.

If you use IMS Image Copy/Recovery utilities to move data and later find that changing block size or reorganization is justified, use the IMS/VS data base reorganization utility.

This approach is helpful for the following reasons:

- Using the IMS/VS Image Copy/Recovery utilities synchronizes the movement of a data base with recovery data and procedures.
- There are limitations on the kinds of recovery that you can perform after a data base has been moved with non-IMS/VS utilities.
- Because the IMS/VS Image Copy/Recovery utilities first make a backup of the data on an intermediate device, then recover to the new device, you have a copy of the migrated data available in case you encounter problems.
- You can back up an IMS/VS data base either in the batch mode, when it is allocated exclusively to the backup process, or in the online mode. Note that if you run Data Base Image Copy concurrently with execution of the online system or another batch system, issue the /DBD command for the data base you are copying and *wait* for message DFS488 *before* starting Data Base Image Copy.
- Backup and recovery procedures are already in place and tested in most installations using IMS/VS. These procedures can probably be used for moving data to 3380s, instead of coding additional JCL or using other tools.
- The selection of optimal block sizes for use with IMS/VS data base data sets is somewhat complex, and must be coordinated with the selection of buffer pool sizes. The designer of a data base has probably analyzed the data access techniques used with specific applications, and may have designed the physical blocking to favor specific applications. Block sizes should **never** be changed without the involvement of the responsible data base administration group. Some specific hazards of changing the block size of IMS/VS data base data sets are:

Large block sizes may adversely affect online performance because many requests may be for small amounts of specific data.

Large block sizes can reduce data concurrency when block level sharing is used.

Some IMS/VS data organizations do not allow the physical location of a segment to be changed in a data base without IMS/VS intervention. These data sets must not be reorganized using non-IMS/VS utilities.

For information about the IMS/VS utilities, see *IMS/VS Version 2 Utilities Reference Manual*.

Moving DB2 Data

You can move DB2 data to all 3380 models from 3350 devices, and you can move DB2 data interchangeably between 3380 models. You can move all DB2 data sets with DFDSS 2.3.0 or later. This section identifies items to consider before moving DB2 data sets, and alternative ways to move them.

Considerations Before Moving DB2 Data Sets

The following items should be considered before moving DB2 data between different DASD devices:

- Data set format. DB2 table spaces and DB2 index spaces consist of one or more VSAM ESDSs. Although DB2 data sets are defined as VSAM ESDSs, they are not formatted like standard VSAM data sets. Specify control interval (CI) mode processing to insure proper handling. Do **not** change the block size of DB2 data.
- Data set management. A user can let DB2 manage the VSAM data sets for table spaces and index spaces by defining DB2 storage groups (a set of DASD volumes that DB2 uses to manage these table space and index space data sets). Alternatively, the user can maintain closer control over the physical storage of tables and indexes by using access method services.
- Data inconsistencies. While not a normal occurrence, data inconsistencies may occur when executing DB2 under IMS/VS or CICS/OS/VS. The contents of a DB2 data set may not be consistent due to unresolved states in DB2 caused when DB2 loses its connection to IMS/VS or CICS/OS/VS and the DB2 data is left uncommitted. Use the DB2 DISPLAY command to determine if there are any unresolved states on the table space and index spaces to be copied. Resolve these inconsistencies prior to the copy with the DB2 RECOVER INDOUBT command or the DB2 RECOVER utility.
- Data integrity. Stop the DB2 table space before moving DB2 data sets contained therein. Issue the DB2 STOP DATABASE command with the SPACENAM parameter or bring down DB2 to flush the DB2 buffers, close the data sets, and prevent the allocation of data sets to DB2. After the data set has been moved, issue the DB2 START DATABASE command or start up DB2.

Moving a DB2 Data Set or an Entire DASD Volume to Another Volume

When moving DB2 data sets between either like or unlike devices within the same DB2 subsystem, it is more efficient to use tools that do a direct copy of the data sets than tools that rebuild the DB2 data sets. DFDSS, DFHSM, DFP EXPORT/IMPORT, and DSN1COPY all do a direct copy of the DB2 data sets. However, using the DB2 COPY/RECOVER utilities has the following advantages:

Using the DB2 Copy/Recovery utilities synchronizes the movement of a data base with recovery data and procedures.

Because the DB2 Copy/Recovery utilities first make a backup of the data on an intermediate device, then recover to the new device, you have a copy of the migrated data available in case you encounter problems.

Backup and recovery procedures are already in place and tested in most installations using DB2. These procedures can probably be used for moving data to 3380s, instead of coding additional JCL or using other tools.

When moving between like devices, a full-volume operation can be done with either a DFDSS COPY or a DFDSS DUMP/RESTORE. A full-volume operation is faster than a data set operation; however, full-volume operations can only be done from one DASD volume to another of equal or greater capacity. The source volume serial number is copied to the target volume when you specify the COPYVOLID parameter. If you want a different serial number on the target volume, then the access method services DEFINE CLUSTER with the RECATALOG parameter must be executed for **each** data set that was moved before it can be accessed on the target volume.

When moving DB2 data sets between unlike devices, do either a DFDSS COPY by data set or a DFDSS DUMP/RESTORE by logical data set. If the target volume has a different serial number, the COPY and RESTORE will automatically update the integrated catalog facility catalog.

If DFDSS is not installed, then DFHSM is the next easiest tool. DFHSM can move one or all the data sets on the source DASD volume with one control statement which migrates the data set(s) and immediately recalls it to the specified target DASD volume. DFHSM can move DB2 data sets between like or unlike devices. To use DFHSM for DB2 data sets, the access method services support for exporting ESDSs by control interval must be specified. If the source and target volumes are not managed by DFHSM, the UNIT parameter must be specified for the source volume and the device type specified for the target volume.

DFHSM can move a DB2 data set to a different DASD volume with the MIGRATE DATASETNAME command with the CONVERT parameter. To migrate an entire DASD volume and immediately recall it to a different DASD volume, use the MIGRATE VOLUME command with the CONVERT parameter. If not authorized to use the DFHSM MIGRATE command, use DFHSM HMIGRATE and HRECALL commands to move a DB2 data set.

If neither DFDSS nor DFHSM is installed, then DFP EXPORT/IMPORT or DB2 DSN1COPY can be used to move DB2 data sets between like or unlike devices.

Moving and Reallocating a DB2 Data Set

When moving a DB2 data set **not** defined using a storage group, and increasing its space allocation at the same time, the following tools can be used:

- DB2 COPY/RECOVER is the easiest to set up. Image copies are taken on a regular basis and the job streams already exist. Also, if you have a very large table space containing multiple data sets that must be reallocated quickly, select DB2 COPY/RECOVER since it processes only the specified data set.
- DB2 DSN1COPY is the fastest tool. It copies directly from the source data set, thus taking the least amount of data transfer time.
- DB2 REORG is neither the fastest or easiest tool to use but it will process all the data sets of a table space.
- DFP EXPORT by Control Interval can also be used.

Considerations Before Moving Storage Groups

If the target volume serial numbers are different from the source volume serial numbers, then the storage group definition must be changed. The SQL ALTER statement is used to change the definition of the storage group.

When adding DASD volumes to a storage group, only add devices that are the same as others in the storage group. However, the individual partitions of a partitioned table space can be independently assigned to separate storage groups.

When moving DB2 table spaces and index spaces defined in DB2 storage groups, remember:

- All volumes belonging to the same DB2 storage group must be of like device types.

- Data sets within a table space can be on different device types only if they belong to different partitions of a partitioned table space.

After moving a DB2 data set from one DASD volume to another and changing the storage group definition to add the new volume, there is no guarantee that the moved table space or index space will remain on the new volume. DB2 will delete the existing data set and then reallocate the data set on the first volume defined for the storage group with sufficient space for the primary space specification if you do any of the following:

- Execute the DB2 REORG utility on the table space or index space

- Execute the DB2 RECOVER utility on the table space or index space

- Recreate the table or index with SQL DROP and SQL CREATE statements.

Moving and Reallocating When Using Storage Groups

The DB2 data set allocation size should not be changed through means other than SQL DROP and SQL CREATE statements when using storage groups because DB2 keeps track internally of the allocation quantity. This allocation quantity is used by DB2 to reallocate the DB2 data sets during DB2 utility operations and when allocating additional data sets when the existing DB2 data set has reached its maximum size. Data Base Migration Aid Utility (**DBMAUI**) facilitates changing the allocation of a fully extended DB2 data set defined using a storage group. DBMAUI simplifies the procedure documented in the *DB2 Operations and Recovery Guide* since DBMAUI has the capabilities to extract the data from the table to enlarge and generate the SQL CREATE, SQL GRANT, and DB2 LOAD statements. For detailed information, refer to the *DBMAUI User's Guide*.

Procedures for Moving DB2 Table Spaces and Indexes

There may be times when you choose to move one or several DB2 data sets, for example: to balance I/O processing across DASD volumes. The following procedures can be used to move DB2 table spaces and indexes between DASD devices within the same DB2 subsystem. The space allocation of the DB2 data set can be changed at the same time it is moved to another DASD volume. Select the procedure that applies to your circumstances.

Moving Table Spaces and Indexes Belonging to a DB2 Storage Group

Note: The sequence of recovering data sets is critical for the DB2 catalog. For the correct sequence, see *IBM DATABASE 2 Operation and Recovery Guide*.

When table spaces and indexes belong to a DB2 storage group, all members of the DB2 storage group must be moved to the new volume. For a DB2 storage group, you should:

1. Make an image copy of each table space with the COPY utility
2. Stop each data base with the -STOP DATABASE command and SPACENAM parameter
3. Delete the old volumes with the ALTER STOGROUP statement
4. Add the new volumes with the ALTER STOGROUP statement
5. Start each data base specifying utility-only access (UT) with the -START DATABASE command and SPACENAM parameter
6. Recover each table space and index with the RECOVER utility
7. Start each data base for read/write access (RW) with the -START DATABASE command and SPACENAM parameter.

Moving Table Spaces and Indexes When You Manage the Data Sets

Note: The sequence of recovering data sets is critical for the DB2 catalog. For the correct sequence, see *IBM DATABASE 2 Operation and Recovery Guide*.

For table spaces and indexes that have many data sets, you can use the following procedure for each table space or index:

1. Make an image copy of the table space with the COPY utility
2. Stop the data base that contains the table space with the -STOP DATABASE command and SPACENAM parameter
3. Delete the original data sets of the table space or index with the access method services DELETE
4. Define the data sets on the target volumes with the access method services DEFINE command
5. Start the data base specifying utility-only access (UT) with the -START DATABASE command and SPACENAM parameter
6. Recover the table space or index with the RECOVER utility
7. Start the data base with read/write access (RW) with the -START DATABASE command and SPACENAM parameter.

For more information about the DB2 commands and utilities referred to in these procedures, see:

IBM DATABASE 2 Reference

IBM DATABASE 2 Operation and Recovery Guide

Chapter 10. Performing Follow-Up Activities

After you have moved the necessary data, there are a number of follow-up and ongoing storage management tasks to perform. These include:

- Tuning the system for performance
- Backing up your data
- Documenting the new configuration and data set placement
- Updating your installation's written operating procedures
- Performing media maintenance

Performing these activities ensures that your 3380s provide acceptable levels of performance. Performance and space analysis helps you identify potential problems before they occur and allows you to plan for orderly storage subsystem growth. Backing up your data on a regular basis, documenting the configuration, and updating operating procedures enables smooth operation of your computing complex and minimizes recovery time in the event of a loss of data. Regular media maintenance maximizes subsystem performance and availability.

This chapter describes these activities and provides references to more detailed information.

Tuning the System for Performance

After you have moved the data to the new volumes, it is a good idea to monitor their activity, along with the activity of the associated channels and storage control. If you find volumes or paths that are overloaded, redistribute the data to relieve the overloaded components.

Use RMF to determine path and device loading as discussed in "Using RMF to Review I/O Workload" on page 22. This will also tell you if you are meeting the performance requirements you have chosen for your storage subsystems. For 3380 performance considerations, see "Performance Considerations" on page 35. For additional information, see *MVS SML: Configuring Storage Subsystems* and other documents listed in the "Bibliography" on page 145.

Use Cache RMF Reporter to aid in tuning the performance of storage subsystems with cached storage controls. Of the many reports it provides, perhaps the most useful are:

Subsystem Status report which shows the cache configuration at the end of the RMF interval or at the end of the specified interval of duration

Device Aggregation report which gives the various hit ratios and read-to-write ratios in addition to the request totals broken down by volume in addition to the control unit totals

Subsystem Summary report which gives the I/O rate, R/W ratio, and the various hit ratios

You can use Cache RMF Reporter to:

- Run summary report for peak period durations:
 - Watch for trends
 - Compare with RMF or other performance data for same periods.
- Run summary report to pick out problematical intervals:
 - Look for spike intervals
 - Compare with RMF or other performance data
 - Try to determine what was going on in the system at that time
 - Determine whether a remedy is required.
- For problematical intervals:
 - Run Subsystem Status report to verify that desired caching configuration is implemented
 - Run device aggregation report to check on staging/destaging
 - Check on RMF pend time and % reserved.
- If RMF response time is too high, look for:
 - Low hit ratios
 - Low read-to-write ratio if not using fast write
 - Connected searches
 - 3380 AA4 internal path contention.

For further information, see *Cache RMF Reporter Program Description/Operations Manual*.

Backing Up Your Data

This topic is discussed in detail in *MVS SML: Managing Data Sets* and *MVS SML: Managing Storage Pools*. The following information provides you with an overview of the task of backing up your data.

In addition to volume or data set backup when migrating data to new devices, you should backup your data at regular intervals to minimize the impact of data loss. Prepare a written backup and recovery procedure for your computer complex. Practice these recovery procedures to minimize outages when recovery is necessary.

Procedures for creating backup copies and strategies for recovery depend both on the frequency of change to the data and on the potential effects if data is temporarily unavailable or not current. Your backup procedure should include incremental and full-volume backups. An incremental backup copies only data sets added or changed since the last backup. Full-volume backups used in conjunction with incremental backups provide a more accurate reconstruction of a volume's contents and make volume recovery faster.

Data Base Backup

For data managed by a data base product (for example, IMS/VS), follow the recommendations given in the product documentation. For IMS/VS, see *IMS/VS Version 2 Operations and Recovery Guide*. For DB2, see *IBM DB2 Operation and Recovery Guide*.

Incremental Backup

DFHSM can perform incremental data set backup on selected volumes automatically. With incremental backup, data sets are backed-up only if there is no backup copy or if the data has changed since the last incremental back up. DFHSM maintains records of backed-up data sets in a control data set, so that you can use DFHSM to recover the data sets with a minimum of effort.

You can use DFDSS for incremental backup by submitting a job that invokes DFDSS and supplies the appropriate commands to initiate the backup operation. Using DFDSS, you can back up all changed data sets on one or more volumes, or you can select the data sets based on a catalog, high-level qualifier, or other criteria. However, you have to maintain a record by retaining the DFDSS output.

Backup and Recovery Considerations

Your backup and recovery plans should consider the following:

- Frequency of backups and number of versions
- Use of DFHSM
- Use of DFDSS
- Use of ISMF
- Data set Recovery
- Volume Recovery
- Deletion of obsolete backup data sets
- Disaster Recovery Plan.

Full-Volume Backup and Recovery

Consider backup of entire volumes at regular intervals, using DFDSS for the following reasons:

- If an entire volume is lost, you can quickly restore the data to a different device.
- All data on the backup volume is synchronized since it is written to the backup at the same time.

The interval at which you backup these volumes depends on how important the data is, how frequently the data changes, and the times at which the data can reasonably be backed up. During a full-volume backup, the entire volume is unavailable to users. This can influence the backup interval; for example, there may only be one time period each week when it is acceptable to prevent user access to the data that needs to be backed up. DFDSS can restore a single data set or a group of data sets from a full-volume backup.

DFHSM Version 2 Release 3.0 (and later releases) invokes DFDSS to perform full volume dumps at specified intervals. This capability enables you to fully automate volume dumps. For more information, see *MVS SML: Managing Data Sets* and *MVS SML: Managing Storage Pools*.

Documenting the New Configuration and Data Set Placement

If you documented your planned configuration as described in “Documenting Configuration and Data Set Placement” on page 57, verify the accuracy of the document. Incorporate any changes made in your system configuration and in volume or data set placement.

In terms of data sets, it is a good idea to keep a record of the following information for the data sets for each of your major applications:

- The name and location of the data set
- The function of the data set
- The performance and security requirements of the data set
- The owner of the data set.

If you haven't documented this information, now is the time to do so. This information is essential for any backup or recovery activity, and is often needed during normal system operation.

Updating Written Operating Procedures

After you have moved the data, you will need to communicate new operating procedures and data set allocation procedures to the user community. In many instances, these procedures are handled automatically, and users need not be informed of changes to them. However, procedures vary from one computer complex to the next. You must determine what new information to communicate to your user community, and the best way to communicate it. In addition, you may want to prepare a brief presentation to supplement or introduce the updated procedures.

Your service level agreements that you have entered into with your user groups may require updating. For more information, see *MVS SML: Managing Data Sets*

Use the information contained in Chapter 7, "Operating the 3380 under MVS" on page 77 to prepare written procedures for your operators. Operators need to know the correct procedure for taking the devices online and offline. They also need to know how the new units are operated (for example, the correct power-on and power-off procedures).

Performing Media Maintenance

Media problems with storage devices affect system performance and hardware availability. Your part in media maintenance includes running the System Exception Reports listed in the following section, analyzing them, separating media problems from hardware problems, and handling the media problems. To find out how to obtain these reports, see *Environmental Record Editing and Printing Program User's Guide and Reference*.

Maintaining IBM Storage Subsystem Media explains in detail how each of these reports contributes to the error handling process. It assists you in interpreting the EREP reports and gives you detailed information for using Device Support Facilities to handle error situations.

You can fix most media problems with Device Support Facilities. The hardware problems are the responsibility of your service representative. For information on using ICKDSF with 3380s, see *Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage*. For detailed information on the Device Support Facilities, see *Device Support Facilities User's Guide and Reference*.

System Exception Reports

The Environmental Record Editing and Printing (EREP) program produces a set of System Exception reports, that are based on the contents of the error recording data set (commonly referred to as LOGREC).

- The **System Error Summary (Part 2)** identifies permanent I/O errors, by job and time, associated with data or equipment checks.
- The **Subsystem Exception DASD** report lists accumulated permanent and temporary I/O errors.
- The **DASD Data Transfer Summary** provides details on data checks.

Run the above reports daily as part of normal processing, and designate a member of the data processing center staff to review these reports for actual or potential disk storage problems.

By monitoring your storage devices, you can minimize disruption from problems. When errors occur and the source has been identified, you can alert the responsible party to take action:

- When the source of the error is hardware, call your service representative for service.
- When the source of the error is the volume (media), use Device Support Facilities to handle the specific error condition.

Using the Service Information Messages

With the 3380 attached to the 3990 Storage Control, or the 3380 Direct Channel Attach Model CJ2, a SIM Alert message is displayed on the operator's console to notify operations personnel that the storage control hardware detected an abnormal condition and has recorded a Service Information Message (SIM) on the Error Recording Data Set (ERDS). It is essential that you run an EREP System Exception Report to get the additional information from the ERDS as quickly as possible. The detailed SIM information includes severity of the abnormal condition and impact of the service action to repair the fault. This information will help you determine when to schedule a service action to repair the fault.

Interpreting the SIM Alert Message

The format for the SIM Alert message is shown in Figure 45.

```
IEA480E cuu,xxxx, yyyy, ALERT, MT= , SER=01aa-dddd, REFCODE=nnnn nnnn nnnn,
```

Figure 45. MVS/370, MVS/XA, and MVS/ESA SIM Alert Message Format

The fields of the MVS SIM Alert messages include:

- xxxx** is the failing component. SCU specifies a storage control fault.
- yyyyyy** is the severity of the failure. The severity can be: ACUTE, SERIOUS, MODERATE, or SERVICE.
- MT =** is the machine type and model number (7 characters maximum).
- SER =** is the serial number of the failing unit. The "01" at the beginning of the SER designation indicates that the unit is an IBM-manufactured device.
- REFCODE =** is a reference code that describes the error and references the list of parts the service representative will need to repair the fault. Thus, the reference code helps the service representative to quickly find the cause of the error and repair the fault.

Figure 46 explains the meaning of the severity field for an SCU failing component.

Severity: SERVICE	Severity: MODERATE	Severity: SERIOUS	Severity: ACUTE
A service-related fault occurred that does not affect storage path operation.	A storage cluster temporary error threshold has been exceeded, but both storage paths are operational.	A permanent error has occurred on one storage path. One storage path remains operational.	A permanent error has occurred on both storage paths in the storage cluster.

Figure 46. Meaning of SIM Alert Severity Field for SCU Failing Component

Establishing SIM Handling Procedures

A SIM alert remains on the system console display until the operator or storage administrator responds to the message. When a SIM alert is displayed, request EREP be run to obtain details of the failure. The EREP control cards for selecting the "A3" records generated by the SIM are described in "Generating EREP Report Requests." For instructions on preparing the required JCL, see *Environmental Record Editing and Printing Program User's Guide and Reference*,

You can automate EREP job submission if you keep the JCL and control cards for the EREP job in one of your procedure libraries. An appropriate command from the operator's console (or TSO terminal functioning as an extended console) can then invoke the EREP job. By using the write-to-operator (WTO) exits in JES2 and JES3, the EREP reports will be generated quickly and automatically should a SIM alert occur. Your role would then consist of examining the output and taking the appropriate action for service.

Generating EREP Report Requests

You can obtain detail information on SIMs by requesting the Detail Edit and Summary Report. In addition, the System Exception Reports summarize SIM activity in the DASD Informational Messages section.

Detail Edit and Summary Reports

The essential parameters for selecting SIM "A3" records are:

DEV=(3380)	Device type of 3380
TYPE=A	Select A3 records
PRINT=(PT)	To request Detail Edit and Summary

When constructing your EREP report request, use parameters according to your installation's conventions. It is recommended that you *not* use DATE and TIME parameters when you request SIM details. If you do, use a date and time period long enough to ensure all SIM messages relating to the SIM alert are reported. Request details as quickly as possible after the alert message is received at the console.

Sample control statements for generating these reports appear in Figure 47. This example requests SIM Detail Edit and Summary reports from the error log (LOGREC) rather than from a history file. For instructions on constructing report requests, see *Environmental Record Editing and Printing Program User's Guide and Reference*.

```
//STEP1 EXEC PGM=IFCEREPI,PARM=('TYPE=A',          Request A3s, keep
          'DEV=(3380)', 'PRINT=PT', 'ACC=N)        records in LOGREC
//SERLOG DD DSN=SYS1.LOGREC,DISP=OLD             Input is LOGREC
//DIRECTWK DD UNIT=SYSDA,SPACE=(CYL,5,,CONTIG)   Workspace
//EREPT1 DD SYSOUT=A,DCB=BLKSIZE=133             Duplicate output
//EREPT2 DD SYSOUT=A,DCB=BLKSIZE=133             for online view
//SYSIN DD DUMMY                                  SYSIN not needed
/*
```

Figure 47. Sample System Exception Report Request

For each selected "A3" record, SIM sense bytes are formatted and evaluated to provide additional details on the error and its effect on system performance. The 32 SIM sense bytes are printed in hex-character format. A sample SIM record from the Detail Edit and Summary report appears in Figure 48.

```
REPORTING DEVICE: 000844 REPORT: ASYNCHRONOUS DAY YEAR
REPORTING DEVICE TYPE: 3380 REPORTING SYSTEM: DOS REL. n DATE: 112 87
REPORTING PATH: 20-0844 HH MM SS.TH
TIME: 23 49 04.76

RECORD TYPE: DASD SIM
DEVICE DEPENDENT DATA
SERVICE INFORMATION MESSAGE
SERIOUS ALERT 3380-CJ S/N 0113-0080602 REFCODE 3121-0000-0004
PERMANENT ERROR(S) ON 1 OF 2 STORAGE PATHS FOR SSID 00CC
REPAIR WILL DISABLE STORAGE PATHS 0 AND 1 FOR 00CC

HEX DUMP OF RECORD
HEADER A3831810 00000000 0087112F 23490476 13021170 30810000
0018 00000000 00000000 00000000 00000000 00000000 20200844 800F2023
0038 08000844 D9C1E2F8 F4F40000 00901000 00008FE0 23800020 00000104 22002782
0058 000426C6 05100212 F1000500
```

Figure 48. Example of a Formatted Detail Edit and Summary Report

If the SIM sense data has some irregular information in it, the EREP program does not format the 32 SIM sense bytes into the usual messages for evaluation. Instead, EREP prints the information as hex characters, with byte numbers specified for readability, as shown in Figure 49 on page 118.

```

REPORTING DEVICE:      0001C4      REPORT: ASYNCHRONOUS      DAY YEAR
REPORTING DEVICE TYPE:  3380      REPORTING SYSTEM: VS n REL. n 370XA DATE: 117 87
REPORTING PATH:       11-01C4      HH MM SS.TH
                                           TIME: 17 01 40.95

RECORD TYPE:          DASD SIM
DEVICE DEPENDENT DATA
DEVICE DEPENDENT DATA NOT FORMATTED
SYSTEM INFORMATION DATA
BYTE 00 01 02 03 04 05 06 07
      D9 C1 E2 F1 F8 F4 00 00

SUBSYSTEM INFORMATION DATA
BYTE 00 01 02 03 04 05 06 07 08 09 10 11 12 13 14 15
      00 90 10 00 00 00 8F E0 23 9C 00 10 00 00 0E 04
BYTE 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31
      22 00 27 7F F5 00 31 21 05 10 42 11 F4 F4 00 00

HEX DUMP OF RECORD
HEADER  A3831810  00000000  0087117F  17014095  13221170  30810000
0018  00000000  00000000  00000000  00000000  00000000  00000000  201101C4  800F2021
0038  080001C4  D9C1E2F1  F8F40000  00901000  00008FE0  239C0010  00000E04  2200277F
0058  F5003121  05104211  F4F40000

```

Figure 49. Example of an Unformatted Detail Edit and Summary Report

System Exception Reports

It is best to run the System Exception Reports as a normal part of your operation. Summary information of SIM activity appears in the DASD Informational Messages report, as shown in Figure 50 on page 119.

The COUNT field on this report indicates the number of times a specific SIM is reported. If duplicate SIMs are received, the message is listed once and the COUNT field indicates how many times it occurred during the time period specified by the report request.

For instructions on requesting the System Exception Reports, see *Environmental Record Editing and Printing Program User's Guide and Reference*.

DASD INFORMATIONAL MESSAGES			REPORT DATE 126 87
			PERIOD FROM 112 87
			TO 117 87
PHYSICAL ID	SYMPTOM CODE	COUNT	MESSAGE

0010111.X	000E	1	SERVICE INFORMATION MESSAGE SERVICE ALERT 3990-02 S/N 0112-0010111 REFCODE 3121-1000-000E PERMANENT ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F500 AND F400 REPAIR WILL DISABLE STORAGE PATHS 0 AND 1 FOR F400 AND F500
0010111.X	0011	1	SERVICE INFORMATION MESSAGE SERVICE ALERT 3990-02 S/N 0112-0010111 REFCODE 2600-3000-0011 PERMANENT ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F500 AND F400 REPAIR WILL DISABLE STORAGE PATHS 2 AND 3 FOR F400 AND F500
0010222.X	0013	1	SERVICE INFORMATION MESSAGE SERVICE ALERT 3990-02 S/N 0112-0010222 REFCODE 2600-2000-0013 TEMPORARY ERROR(S) ON 1 OF 4 STORAGE PATHS FOR SSID F900 REPAIR WILL DISABLE STORAGE PATHS 0 AND 1 FOR F900
0080602.X	N/A	1	SERVICE INFORMATION MESSAGE SERVICE ALERT 3380-CJ S/N 0113-0080602 REFCODE 3121-0000-0004 PERMANENT ERROR(S) ON 1 OF 2 STORAGE PATHS FOR SSID 00CC REPAIR WILL DISABLE STORAGE PATHS 0 AND 1 FOR 00CC
0080608.X	N/A	1	SERVICE INFORMATION MESSAGE SERVICE ALERT 3380-CJ S/N 0113-0080608 REFCODE 2818-0000-0000 TEMPORARY ERROR(S) ON 1 OF 2 STORAGE PATHS FOR SSID 00EE REPAIR WILL DISABLE STORAGE PATHS 0 AND 1 FOR 00EE

Figure 50. Example of SIMs in System Exception Informational Messages

Appendix A. Sample DASD Installation and Migration Project Plan

This appendix presents a sample project plan designed for a typical MVS DASD installation and migration scenario.

The project plan documents the following:

- The tasks that need to be completed, including interdependences among tasks
- The implementation schedule
- Actual completion dates for tasks
- Responsible individuals

You can produce adequate planning documents and schedules with simple tools. You can prepare planning information similar to that on the following pages using an editor that has sort-on-column capability. Then a variety of reports can be generated. For example, reports for each person who has responsibility for tasks, lists in planned-completion-date sequence, lists of behind-schedule tasks, and so forth. Useful project plans can also be manually generated and maintained. No matter how you create them, your plans should be written down and understood by all involved.

The sample project plan on the following pages groups tasks by area of activity. Whether you are replacing all of your old DASD with new DASD, or adding a new string, document and review your plans to make your transition to the new hardware as smooth as possible.

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> INSTALLATION PLANNING

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
A001	Identify requirements for floor space, power air conditioning, cables						
A002	Identify prerequisite features						
A003	Verify iogen requirements						
A004	Verify equipment delivery schedules						
A005	Identify licensed program requirements: order, install and test programs						
A006	Identify and install software tools						
A007	Review vendor and application software support						
A008	Identify and apply software fixes						
A009	Identify publication requirements						
A010	Create project team						
A011	Assign responsibility for project plan						
A012	Schedule system time for volume preparation and data movement						
A013	Plan for personnel training						
A014	Document new operating and programming procedures						

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> UNDERSTANDING CURRENT DATA AND HARDWARE CONFIGURATION

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
B001	Inventory data sets						
B002	Identify performance-critical data sets						
B003	Identify availability and space requirements						
B004	Identify device dependencies						
B005	Measure effectiveness of current hardware configuration						
Notes:							

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> PLANNING HARDWARE CONFIGURATION

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
C001	Understand hardware functions and capabilities						
C002	Understand capacity considerations						
C003	Understand performance considerations						
C004	Understand availability considerations						
C005	Understand shared DASD considerations						
C006	Understand guest system considerations						
C007	Plan configurable units						
C008	Plan I/O addresses						
C009	Document planned hardware configuration						
Notes:							

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> PLANNING DATA CONFIGURATION

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
D001	Identify performance-critical data						
D002	Identify candidates for cache volumes						
D003	Plan data movement strategy by volume and/or data set						
D004	Plan for backup and recovery during data movement						
D005	Document data configuration and data movement strategy						
D006	Conduct a review of the data move strategy						
D007	Schedule system time for data movement						
Notes:							

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> INSTALLING DASD

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
E001	Assign I/O addresses to new devices						
E002	Define devices to system with either iogen, sysgen, or MVSCP						
E003	Define devices to processor with IOCP generation						
E004	Physically install new units						
E005	Calculate size of VTOC and VTOC index						
E006	Initialize volumes with ICKDSF						

Notes:

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> PLANNING OPERATIONS CHANGES

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
F001	Document changes in operator control panels						
F002	Document power on/off procedures for devices and storage controls						
F003	Document enable/disable procedures for devices and storage controls						
F004	Document procedures to display device status						
F005	Document procedure to vary online/offline devices and storage controls						
F006	Document procedure for nondisruptive DASD installation						
F007	Document SIM alert procedures						
F008	Train operators in new procedures (service and/or remote support)						
Notes:							

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> MOVING DATA

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
G001	Identify tools for moving data						
G002	Back up volumes before data migration						
G003	Back up data sets before data migration						
G004	Move unmovable data sets						
G005	Move full volumes						
G006	Move system data sets						
G007	Move IMS/VS data sets						
G008	Move DB2 data sets						
G009	Move CICS data sets						
G010	Move VSAM data sets						
G011	Move non-VSAM data sets						
G012	Move other application data sets						
G013	Move tape data sets						
Notes:							

DASD INSTALLATION AND MIGRATION PROJECT PLAN

FOCUS AREA ==> MONITORING AND MAINTAINING DEVICES

TASK ID	TASK DESCRIPTION	DEPENDENCIES	PLANNED		ACTUAL		PERSON RESPONSIBLE
			START	END	START	END	
H001	Monitor system for SIM alerts						
H002	Initialize RMF to collect performance statistics						
H003	Tune storage subsystem for performance						
H004	Run EREP system exception reports to track DASD errors						
H005	Use ICKDSF to fix media errors						
Notes:							

Acronym List

This list contains definitions for acronyms and abbreviations used in the various books in the Storage Subsystem Library. The terms in this list are not necessarily used in *this specific* book. Some terms are more specifically defined in the glossary.

CCHH	cylinder, cylinder, head, head	ID	identifier
CCHHR	cylinder, cylinder, head, head, record	IML	initial microcode load
CCW	channel command word	IOCP	I/O configuration program
CHL-I	channel interface	I/O	input/output
CHPID	channel path identifier	IPL	initial program load
CI	control interval	ISMF	Interactive Storage Management Facility
CKD	count-key data	ISPF/PDF	Interactive System Productivity Facility / Program Development Facility
CMS	Conversational Monitor System	JCL	job control language
CP	control program	KL	key length
CTL-I	control interface	KVA	kilovolt x ampere
DASD	direct access storage device	LRU	least recently used algorithm
DCB	data control block	Mb	megabyte
DFDSS	Data Facility Data Set Services	MAP	maintenance analysis procedure
DFHSM	Data Facility Hierarchical Storage Manager	NVS	nonvolatile storage
DFSORT	Data Facility SORT	R0	record zero
DL	data length	RACF	Resource Access Control Facility
DLS	device level selection	RMF	Resource Measurement Facility
DLSE	device level selection enhanced	RPS	rotational position sensing
DPS	dynamic path selection	RTM/SF	Realtime Monitor / Systems Facility
ECKD	extended count-key data	SCA	shared control array
EOF	end-of-file	SF	support facility
ERDS	error recording data set	SIM	service information message
EREP	Environmental Record Editing and Printing	SLR	Service Level Reporter
FBA	fixed-block architecture	SMF	System Management Facilities
FCCHH	Flag, cylinder, cylinder, head, head	SMS	Storage Management Subsystem
FRU	field replaceable unit	SQL/DS	Structured Query Language / Data System
Gb	gigabyte	SSID	subsystem identifier
GBOF	general bill of form	TCO	Triple Capacity Option
GRS	MVS global resource serialization	TPF	Transaction Processing Facility
HA	home address	UCB	unit control block
HDA	head-disk assembly	VMMAP	VM Monitor Analysis Program
		VMPPF	VM Performance Planning Facility
		VTOC	volume table of contents

Glossary

This glossary contains disk storage subsystem terms used in the various books of the Storage Subsystem Library (SSL). To help explain some of the terms related to configuration of storage subsystems, several illustrations are included at the end of this glossary. The definitions of certain terms include references to these illustrations. Some of the illustrations may reflect hardware configurations not supported in your operating environment.

Each of the terms included here is not necessarily used in *this specific* book. If you do not find the term you are looking for, refer to the index or to *Dictionary of Computing*, SC20-1699.

A-unit. The direct access storage unit that contains the controller functions to attach to the storage control. An A-unit controls the B-units that are attached to it and is often referred to as a head of string.

access authorization. Bits in the Define Extent file mask that define one of the three authorization groups for a channel program (normal authorization, device support authorization, or diagnostic authorization).

access mechanism. See actuator.

active duplex state. A state of operation that occurs when both devices in a dual-copy logical volume are automatically updated. See also duplex state and suspended duplex state.

actuator. A set of access arms and their attached read/write heads, which move as an independent component within a head and disk assembly (HDA). For example, the 3380 Model AK4 has two HDAs, each containing two actuators. See also device and volume.

alternate track. On a direct access storage device, a track designated to contain data in place of a defective primary track.

B-unit. A direct access storage unit that attaches to the subsystem through an A-unit. A B-unit has no controller functions.

C-unit. A direct channel attach 3380 direct access storage unit that contains both the storage control functions and the DASD controller functions. A 3380 C-unit functions as a head of string and controls the B-units that are attached to it.

cache. A random access electronic storage in selected storage controls used to retain frequently used data for faster access by the channel. For example, 3880 Model 23 and 3990 Model 3 contain cache.

cache fast write. A form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging. This 3990 Model 3 Storage Control function should be used for data of a temporary nature, or data which is readily recreated, such as the sort work files created by the appropriate release of DFSORT.

cache fast write data. Data that the channel command modifies in cache and not on DASD. It has read-hit performance benefits for write hits. See cache fast write.

central complex. All the subsystems in a storage control.

channel connection address. The I/O address that uniquely identifies an I/O device to the channel during an I/O operation.

channel interface (CHL-I). The circuitry of a storage control that attaches storage paths to a host channel.

check-1 error. In the storage control and DASD, an error that does not allow the use of normal machine functions to report details of the error condition.

check-2 error. In the storage control and DASD, an error that can be reported using the normal machine functions.

cluster. See storage cluster.

concurrent maintenance. A 3990 Model 2 and 3 capability that permits a service representative to perform a service action on one storage cluster while normal DASD access operations continue on the other cluster.

On the 3990 Model 3, a service representative can perform most service actions on *nonvolatile storage* while caching and DASD access operations continue through both the storage clusters.

On the 3990 Model 3, a service representative can perform most service actions on *cache* while DASD access operations continue through both the storage clusters. During a cache service action, cache fast write and DASD fast write operations are not performed and dual copy operations are performed directly with the DASD.

connection check alert. The electronic signal used by the 3380 to indicate a check-1 error condition to the storage control. See check-1 error.

contingent allegiance. A state the storage path establishes for an I/O device that allows only the channel path that is communicating with the I/O device to continue to do so. This state occurs when the channel accepts a status byte that contains unit check.

control interface (CTL-I). The hardware connection between the storage control function and the DASD controller function.

control unit. A hardware device that controls the reading, writing, or displaying of data at one or more input/output devices. See also storage control.

controller. The hardware component of a DASD head of string unit that provides the path control and data transfer functions. For example, there are two controllers in a 3380 Model AE4, AK4, or CJ2.

controller address. The 1-bit address used by the storage control to direct commands to the correct DASD string on the CTL-I. Controller address applies to the 3380 Models AA4, AD4, and AE4. See also string address.

controller ID. An 8-bit identifier that uniquely identifies the physical string regardless of the selection address. It identifies to the service representative, by means of EREP, a failing subsystem component (controller or device) without having to translate a selection address (which may have little relation to a physical address) to a physical component. The controller ID is the number shown on the operator panel. Controller ID applies to the 3380 Models AA4, AD4, and AE4. See also string ID.

count-key-data (CKD). A DASD data recording format employing self-defining record formats in which each record is represented by a count area that identifies the record and specifies its format, an optional key area that may be used to identify the data area contents, and a data area that contains the user data for the record. CKD is also used to refer to a set of channel commands that are accepted by a device that employs the CKD recording format. See extended count-key-data architecture.

DASD. Direct access storage device; for example, a 3380.

DASD fast write. A form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to the DASD. Both copies are retained in the storage control until the data is completely written to the DASD, providing data integrity equivalent to writing directly to the DASD. DASD fast write is available with a 3990 Model 3 Storage Control.

DASD subsystem. One or more DASD strings and the storage control(s) to which the DASD are attached.

demotion. The process of removing the image of one or more records from cache. A set of one or more DASD records is demoted either by being selected for replacement (overlay) by another set of DASD records or by being marked invalid. Compare to promotion.

destage. The asynchronous write of new or updated data from cache or nonvolatile storage to DASD. This is used only for the fast write and dual copy functions of 3990 Model 3. See also fast write and write hit.

device. A uniquely addressable part of a DASD unit that consists of a set of access arms, the associated disk surfaces, and the electronic circuitry required to locate, read, and write data. See also volume.

device address. Three or four hexadecimal digits that uniquely define a physical I/O device on a channel path in System/370 mode. The one or two leftmost digits are the address of the channel to which the device is attached. The two rightmost digits represent the unit address.

device ID. An 8-bit identifier that uniquely identifies a physical I/O device.

device level selection (DLS). A DASD function available with 3380 Models AD4, BD4, AE4, BE4, AJ4, BJ4, AK4, BK4, and CJ2. With DLS, each of the two controllers in the DASD string has a path to all devices in the string (as many as 14 addresses for a CJ2 or 16 addresses for other string types), and any two devices in the 2-path DASD string can read or write data simultaneously. See DLS support mode, and see Figure 51 on page 141 and Figure 52 on page 142.

device level selection enhanced (DLSE). A DASD function available with 3380 Models AJ4, BJ4, AK4, and BK4. With DLSE, each of the four controllers in the 4-path DASD string (as a result of interconnecting two A-units), has a path to all devices in the string (as many as 32 addresses), and any four devices in the 4-path DASD string can read or write data simultaneously.

See DLSE support mode, and see Figure 53 on page 143 and Figure 54 on page 144.

device number. Four hexadecimal digits that logically identify an I/O device in a System/370 Extended Architecture or Enterprise Systems Architecture/370 Systems.

device release. A command that terminates the reservation of the device from the channel issuing the command or from all channels on the interface path group.

device reserve. A command that reserves the device for the channel issuing the command, or for all channels in the same interface path group.

device support authorization. Channel programs executing with this authorization can access all tracks in all track groups, and can execute all Locate Record operations.

Device Support Facilities program (ICKDSF). A program used to initialize DASD at installation and provide media maintenance.

device support tracks. Reserved tracks of a DASD volume that store defect skipping information. This information is used by host utility programs such as ICKDSF.

diagnostic authorization. Channel programs using diagnostic authorization can access the diagnostic and device support tracks only.

diagnostic tracks. Tracks used by the diagnostic programs for testing the read/write function.

director. See storage director.

director-to-device connection (DDC). The control interface that connects a storage path in the storage control to a controller in the DASD A-unit.

diskette drive. A direct access storage device that uses diskettes as the storage medium. A 3880 uses a read-only diskette drive for microcode storage; a 3990 and a 3380 Model CJ2 use a read/write diskette drive for microcode storage and storage control error logs.

DLS support mode. A mode of operation in a 3990 Storage Control that supports 3380 2-path strings, including 3380 AA4 strings and 3380 AD4, AE4, AJ4, and AK4 2-path strings. DLS support mode must be specified by the IBM service representative at installation for the 3990. See single-path storage director, and see Figure 52 on page 142.

DLSE support mode. A mode of operation in a 3990 Model 2 or 3 Storage Control that supports 3380 AJ4 and AK4 4-path strings. DLSE support mode must be specified by the IBM service representative at

installation time for the 3990. See multipath storage director, and see Figure 53 on page 143 and Figure 54 on page 144.

domain. A scope of operations control that spans all the parameters specified by the Locate Record command. For example, the Locate Record command allows only certain commands, and the commands must be in a correct sequence.

See also locate record domain.

DPS array. An electronic storage area that contains device status information. When any 3380 A-unit, except Model A04, is attached to a 3880 Storage Control, the DPS array resides in the A-units. When the same models are attached to a 3990, the DPS array function is part of the 3990 shared control array. The 3380 Model CJ2 contains DPS array in the storage control function.

dual copy. A high availability function made possible by nonvolatile storage in a 3990 Model 3. Dual copy maintains two functionally identical copies of designated DASD volumes in the logical 3990 Model 3 subsystem, and automatically updates both copies every time a write operation is issued to the dual-copy logical volume.

dual-copy logical volume. A logical volume comprised of two physical devices with all data recorded twice, once on each device. A 3990 Model 3 Storage Control automatically ensures that both devices are updated with each write operation to the dual-copy volume. Also called a duplex pair.

dual-frame configuration. Consists of two like storage controls physically interconnected. Pairs of 3880 Model 13 or Model 23 and 3990 Model 2 or Model 3 Storage Controls can be dual-framed. In a dual-frame configuration, each storage director in a logical DASD subsystem is in a different storage control. When a 3990 Storage Control is in DLS support mode, each DASD string has one path to a single-path storage director in each of the 3990 Storage Controls. When a 3990 Storage Control is in DLSE support mode, each DASD string has two paths to a multipath storage director in each of the 3990 Storage Controls.

duplex pair. See dual-copy logical volume.

duplex state. Two devices in a 3990 Model 3 subsystem are in duplex state when they have been made into a dual-copy logical volume.

dynamic path reconnect. A function of dynamic path selection (DPS) that allows disconnected DASD operations to reconnect over any available channel path rather than being limited to the one on which the I/O operation was started. It is available only on System/370 Extended Architecture and Enterprise Systems Architecture/370 Systems. For example, when

a 3990 Storage Control (in DLSE support mode) having four host channels is connected to a 3380 Model AJ4 or AK4 4-path string, any device can reconnect on any one of four completely independent data paths, providing improved performance and availability.

dynamic path selection (DPS). DASD subsystem functions available with all 3380 heads of string except Model A04. These functions include:

- Two controllers providing data paths from the 3380 strings to the storage directors
- Simultaneous transfer of data over two paths to two devices, providing the two devices are on separate internal paths within the string
- Sharing DASD volumes by using System-Related Reserve and Release
- Providing dynamic path reconnect to the first available path.

Environmental Record Editing and Printing (EREP)

program. The program that formats and prepares reports from the data contained in the Error Recording Data Set (ERDS).

erase. To remove data from a data medium, leaving the medium available for recording new data.

error burst. A sequence of bit errors counted as one unit, or burst.

error correcting code (ECC). A code designed to detect and correct error bursts by the use of check bytes.

extended count-key-data (ECKD) architecture. A set of channel commands that use the CKD track format. This architecture employs the Define Extent and Locate Record commands to describe the nature and scope of a data transfer operation to the storage control to optimize the data transfer operation. The 3990 Storage Control supports the ECKD architecture.

extent. A set of consecutively addressed tracks that a channel program can access. The limits of an extent are defined by specifying the addresses of the first and last tracks in the extent.

fast dual copy. A dual copy capability where DASD fast write and dual copy are active concurrently to provide a significant dual copy performance enhancement.

fast write. In a 3990 Model 3 Storage Control, a write operation at cache speed that does not require immediate transfer of data to a DASD. The data is written directly to cache and/or nonvolatile storage and is available for later destaging. Fast write reduces the

time an application must wait for the I/O operation to complete. See also DASD fast write, cache fast write, and destage.

fence. To separate one or more paths or elements from the remainder of the logical DASD subsystem. The separation is by logical boundaries rather than power boundaries. This separation allows isolation of failing components so that they do not affect normal operations.

fixed-block architecture (FBA). A DASD recording format employing data blocks of fixed size. The blocks are addressed by block number relative to the beginning of the particular file. Contrast with count-key-data (CKD).

gigabyte (Gb). 10^9 bytes.

head and disk assembly (HDA). A field replaceable unit in a direct access storage device containing the disks and actuators. A 3380 Model AK4 has two HDAs.

head of string. The unit in a DASD string that contains controller functions. For example, a 3380 Model AE4, AK4, or CJ2.

home address (HA). The first field on a CKD track that identifies the track and defines its operational status. The home address is written after the index point on each track.

ICKDSF. See Device Support Facilities program.

IDCAMS. A component of Data Facility Product that is also referred to as access method services.

identifier (ID). A sequence of bits or characters that identifies a program, device, controller or system.

IML device. The diskette drive that reads the microcode for the storage control. See diskette drive.

index point. The reference point on a disk surface that determines the start of a track.

initial microcode load (IML). The act of loading microcode.

invalidation. The process of removing records from cache because of a change in status of a subsystem facility or function, or because of an error while processing the cache image of the set of records.

When such a cache image is invalidated, the corresponding records cannot be accessed in cache and the assigned cache space is available for allocation.

I/O device. An addressable input/output unit, such as a direct access storage device, magnetic tape device, or printer.

kilobyte (Kb). 1024 bytes.

least recently used algorithm (LRU). The algorithm used to identify and make available the cache space that contains the least recently used data.

locate record domain. The part of a channel command chain immediately following a Locate Record command that must follow the restrictions imposed by operations the Locate Record parameters specify. The domain is in effect for the number of records or tracks that the count parameter specifies.

See also domain.

logical DASD subsystem. Two storage directors attached to the same DASD strings together with those DASD strings. See Figure 51 on page 141, Figure 53 on page 143, and Figure 54 on page 144.

maintenance analysis procedure (MAP). A step-by-step procedure for tracing a symptom to the cause of a failure.

media. The disk surface on which data is stored.

megabyte (Mb). 10⁶ bytes.

multipath storage director. A storage director in a 3990 Storage Control operating in DLSE support mode. Each multipath storage director in a storage control is associated with two storage paths. All storage paths in a multipath storage director respond to the same range of control unit addresses on a channel. See Figure 53 on page 143 and Figure 54 on page 144.

multitrack operations. A mode of operation in which the storage director advances to the next track when the operation continues past the end of a track.

nondisruptive install. Provides for the physical installation of additional Enhanced Subsystem B-units to an existing 4-path DASD string or an additional 4-path DASD string, concurrently with customer operations, providing access to existing data when DASD unit installation activity is occurring. Nondisruptive install uses the quiesce path and resume path functions and is available when only 4-path Enhanced Subsystem DASD are attached to a 3990 Model 2 or Model 3 Storage Control.

nonvolatile storage (NVS). Additional random access electronic storage with a backup battery power source, available with a 3990 Model 3 Storage Control, used to retain data during a power failure. Nonvolatile storage, accessible from all storage directors, stores data during DASD fast write and dual-copy operations.

normal authorization. A channel program executing with normal authorization cannot access the diagnostic or device support tracks.

orient. An operational code of the Locate Record command that prepares the storage director to position the DASD to the seek address and sector number parameters.

orientation. A control state within a storage path that indicates the type of area (home address, count, key, or data field) that has just passed under the read/write head of the device.

physical ID. A unique designation to identify specific components in a data processing complex.

pinned data. Data that is held in a 3990 Model 3 Storage Control, because of a permanent error condition, until it can be destaged to DASD or until it is explicitly discarded by a host command. Pinned data exists only when using fast write or dual-copy functions.

predictable write. A fast write operation that formats, in cache only, the entire user area of the track and creates a track image. This full-track image is available for later destaging to a DASD.

primary device. One device of a dual-copy volume. All channel commands to the dual-copy logical volume are directed to the primary device. The data on the primary device is duplicated on the secondary device. See also secondary device.

primary track. On a direct access storage device, the original track on which data is stored. See also alternate track.

promotion. The process of moving a track image from a DASD to cache.

quiesce storage path. A function on a 3990 Model 2 or 3 Storage Control in DLSE support mode, configured with only 4-path strings. This function makes one storage path of a multipath storage director unavailable to the processor while assuring that the other path is available for data transfer. This function is initiated by a service representative. Contrast with resume storage path.

read hit. When data requested by the read operation are in the cache.

read miss. When data requested by the read operation are not in the cache.

release. A facility that allows other host systems to communicate with the reserved device. Contrast with reserve.

reserve. A facility for devices attached to multiple channel paths. It allows only one host system to communicate with the specified device. Contrast with release.

resume storage path. A function on a 3990 Model 2 or 3 Storage Control in DLSE support mode, configured with only 4-path strings. This function enables a storage path that has been quiesced. This function is initiated by a service representative. Contrast with quiesce storage path.

rotational position sensing (RPS). A function that permits a DASD to reconnect to a block multiplexer channel when a specified sector has been reached. This allows the channel to service other devices on the channel during positional delay.

secondary device. One of the devices in a dual-copy logical volume that contains a duplicate of the data on the primary device. Unlike the primary device, a limited subset of channel commands may be directed to the secondary device. See also primary device.

service information message (SIM). A message, generated by the host processor upon receipt of sense information from a 3990 or a 3380 Model CJ2, that

contains notification of a need for repair or customer action. The SIM identifies the affected area of the storage control and the effect of the expected service action. A host Error Recovery Procedure (ERP) causes a SIM Alert to be sent to the operator console.

shared control array (SCA). An electronic storage area in a 3990 Storage Control containing the DPS array function and other status information. The shared control array contains status information about its own cluster and the other storage clusters in the storage control subsystem. The information contained in the shared control array is replicated in each storage cluster in that subsystem.

SIM Alert. An operator console message that alerts the operator that an action requiring attention has occurred. The service information message (SIM) can be obtained from the EREP exception report.

simplex state. A volume is in the simplex state if it is not part of a dual-copy logical volume. Terminating a dual-copy logical volume returns the two devices to the simplex state. In this case, there is no longer any capability for either automatic updates of the secondary device or for logging changes, as would be the case in suspended duplex state.

single-frame configuration. In a single-frame configuration, the storage directors of a logical DASD subsystem are located inside one storage control.

single-path storage director. A storage director in a 3990 or 3380 Model CJ2 operating in DLS support mode. Each single-path storage director in the storage cluster is associated with one storage path. A storage path on a single-path storage director responds to a unique control unit address on the channel. A single-path storage director in a 3990 is like a storage director in a 3880. See Figure 52 on page 142.

stage. The process of writing data from a DASD to the cache.

state-change interruption. A combination of bits in the status byte that occurs for a change in the subsystem or the device. For example, a state-change interruption can occur when a volume changes from simplex to duplex. The bit combination is attention, device end, and unit exception. This interruption is sent to all hosts to inform them of the state change.

This was formerly called pack-change interruption.

storage cluster. In the 3990 Storage Control and 3380 Model CJ2, a power and service region containing two independent transfer paths and either one multipath storage director or two single-path storage directors. It is designed so that should a failure or maintenance action occur, it will be independent of the other storage cluster in a 3990 Model 2 or Model 3 Storage Control. The 3990 Model 1 and the 3380 Model CJ2 each have a

single storage cluster; the 3990 Model 2 and Model 3 each have two storage clusters.

In the 3990 Model 3, cache and nonvolatile storage are shared by the storage paths, but are logically and physically separate from the storage clusters. See also storage director, single-path storage director, and multipath storage director.

storage control. The component in a DASD subsystem that connects the DASD to the host channels. It performs channel commands and controls the DASD devices. For example, the 3990 Model 2 and Model 3 are storage controls.

storage director. In a 3990 storage control, a logical entity consisting of one or more physical storage paths in the same storage cluster. In a 3880, a storage director is equivalent to a storage path. See also storage path, single-path storage director, and multipath storage director.

storage director ID. For 3880 Storage Control configurations, an 8-bit designation that uniquely identifies the storage director regardless of its selection address. It identifies to the service representative, by means of EREP, a failing subsystem component (storage director) without having to translate a selection address (which may have little relation to a physical address) to a physical component. The storage director ID is the number shown on the operator panels of 3880s and the attached DASD units.

storage facility. See 4-path string.

storage management subsystem (SMS). An operating environment that helps automate and centralize the management of storage. To manage storage, SMS provides the storage administrator with control over data class, storage class, management class, storage group and automatic class selection routine definitions.

storage path. The hardware within the 3990 Storage Control that transfers data between the DASD and a channel. See also storage director.

storage subsystem. One or more storage controls and their attached storage devices.

string. A series of connected DASD units sharing one or more controllers (or heads of string). For example, a 3380 Model AE4 with the attached B-units is one string.

string address. The 1-bit address used by the storage control to direct commands to the correct 3380 AJ4 and/or AK4 DASD string on the CTL-I. See also controller address.

string ID. An 8-bit identifier that uniquely identifies the physical string regardless of the selection address. It

identifies to the service representative, by means of EREP, a failing subsystem component (controller or device) without having to translate a selection address (which may have little relation to a physical address) to a physical component. The string ID is the number shown on the operator panel of the 3380 Model AJ4 or AK4. See also controller ID.

substring. In a 4-path Enhanced Subsystem DASD configuration, one of the two A-units and the physically adjacent B-units (as many as three B-units).

subsystem identifier (SSID). In a 3990 Storage Control configuration, a number that identifies the physical components of a logical DASD subsystem. This number is set by the service representative at time of installation, and is included in the vital product data in the support facility. This number is identified on the 3380 Enhanced Subsystem models and 3990 operator panels.

subsystem storage. A term used for cache in a 3880 Model 13 or 23. See cache.

support facility (SF). A component of each 3990 and 3380 Model CJ2 storage cluster that provides initial microcode load, error logging, maintenance panel, MAPs, and microdiagnostic functions for that cluster.

suspended duplex state. When only one of the devices in a dual-copy logical volume is being updated because of either a permanent error condition or an authorized user command. All writes to the remaining functional device are logged. This allows for automatic resynchronization of both volumes when the dual-copy logical volume is reset to the active duplex state.

system-managed storage. An approach to storage management in which the system determines data placement and an automatic data manager handles data backup, movement, space and security.

unit address. The last two hexadecimal digits of a DAS device address. This identifies the storage control and DAS string, controller, and device to the channel subsystem. Often used interchangeably with control unit address and device address in System/370 mode.

vital product data (VPD). Nonvolatile data that includes configuration data, machine serial number, EC level, and machine features. It is maintained by the 3990 support facility.

volume. The DASD space accessible by a single actuator. A 3380 Model AK4 contains four volumes, each with 1.89 gigabytes of space.

write hit. When data requested by the write operation are in the cache.

write miss. When data requested by the write operation are not in the cache.

2-path string. A series of physically connected DASD units in which the head of string unit provides two data transfer paths that can operate simultaneously.

4-path string. A series of physically connected DASD units in which the heads of string provide four data transfer paths that can operate simultaneously. A 4-path string requires two 3380 Enhanced Subsystem model A-units.

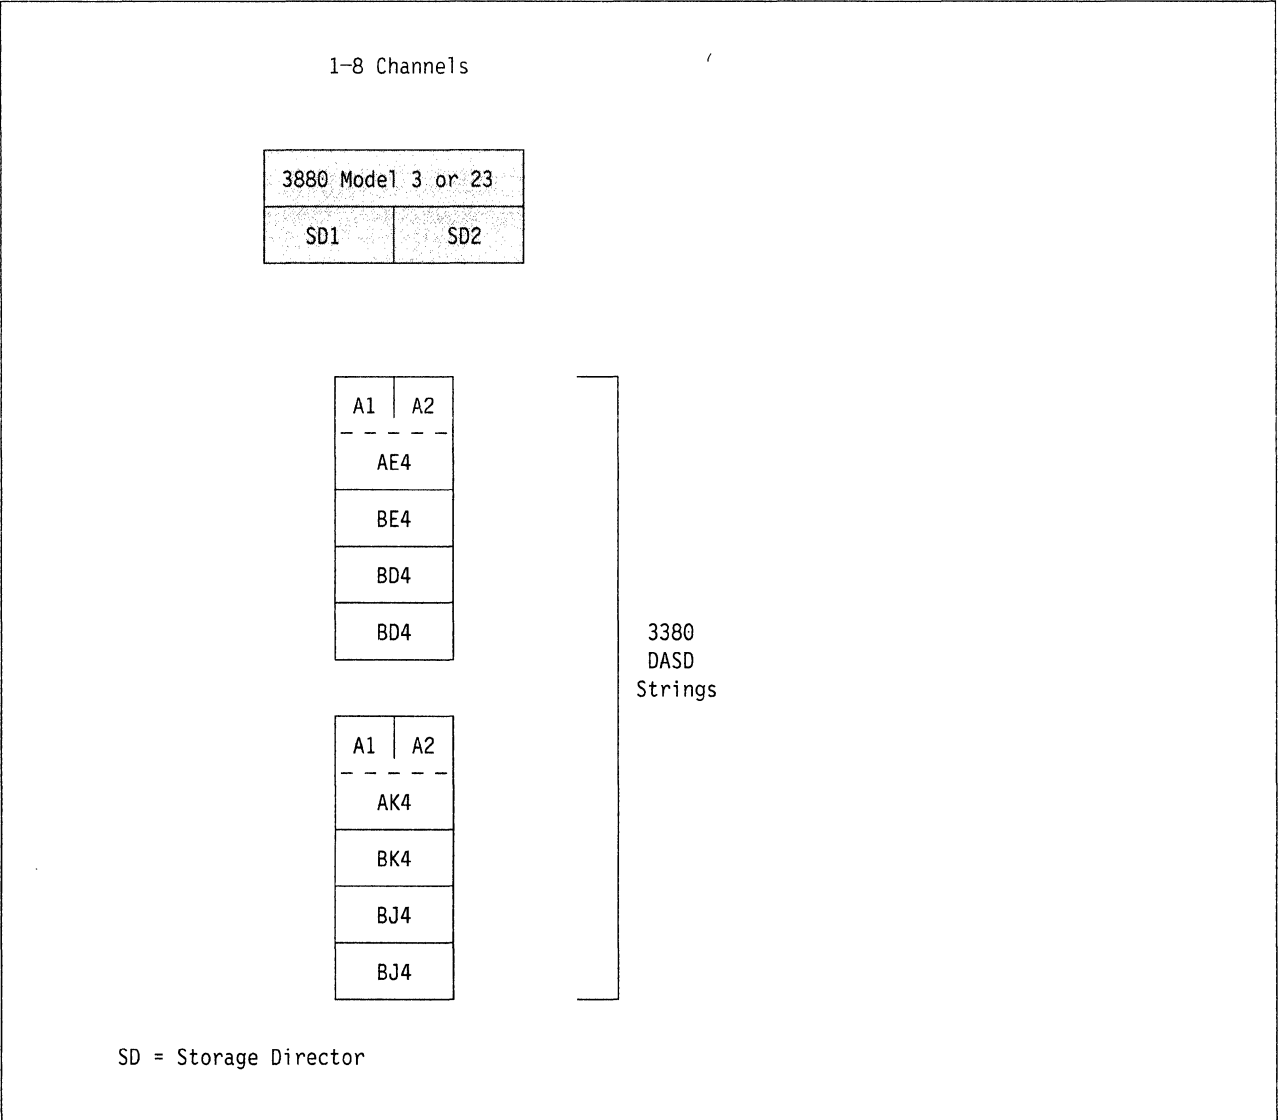


Figure 51. Example of 3380 Model AE4 and AK4 Two-Path Strings Attached to a 3880. Two 3380 strings sequentially connected to both storage directors of a 3880 Model 3 or 23 with appropriate MESS installed. In this example, upper string contains a 3380 Model AE4 controller and a mixture of BD4 and BE4 units. The lower string contains a Model AK4 controller, and a mixture of BJ4 and BK4 units. In a 3880, a storage director performs the same functions as a storage path.

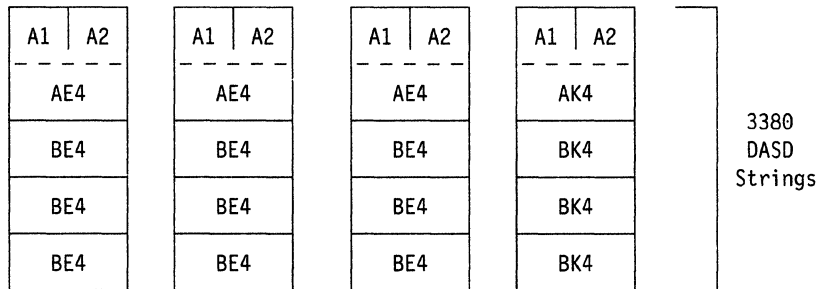

This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLS
- Logical DASD subsystem

As Many as 8 Channels

As Many as 8 Channels



SPSD = Single-Path Storage Director
SP = Storage Path

Figure 52. Example of 3380 Two-Path Strings Attached to a 3990 Model 2 or 3 in DLS Support Mode. This example shows two logical DASD subsystems. Storage Paths 0 and 2 and the attached DASD comprise one subsystem, with a unique subsystem ID, while Storage Paths 1 and 3 and the attached DASD comprise a second subsystem with a unique subsystem ID. Compare this 2-path Enhanced Subsystem string with 4-path Enhanced Subsystem strings shown in Figure 53.

Definitions of the following glossary terms refer to this illustration:

- DLS
- DLS support mode
- Single-path storage director

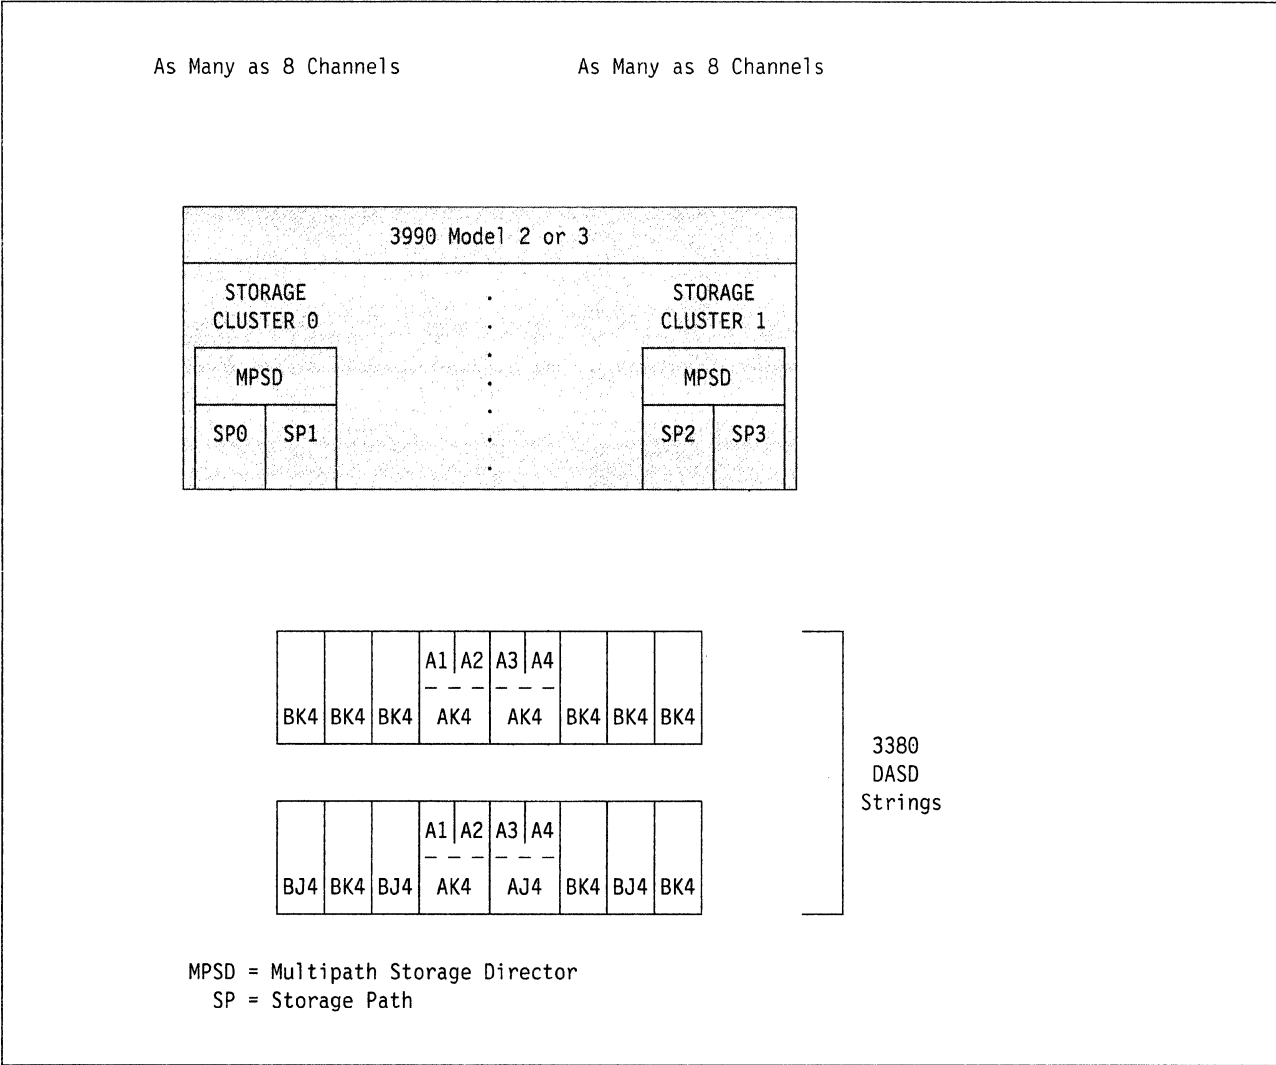


Figure 53. Example of 3380 Enhanced Subsystem Model Four-Path Strings Attached to a 3990 in DLSE Support Mode. Two 3380 4-path strings sequentially connected to the multipath storage directors in the same 3990 Model 2 or Model 3.




This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLSE
- DLSE support mode
- Logical DASD subsystem
- Multipath storage director

As Many as 8 Channels

As Many as 8 Channels



MPSD = Multipath Storage Director
SP = Storage Path

Figure 54. Example of 3380 Two-Path and Four-Path Strings Attached to a 3990 in DLSE Support Mode. Two 3380 2-path strings and one 4-path string, sequentially connected to the multipath storage directors in the same 3990 Model 2 or 3.

This example shows one logical DASD subsystem.

Definitions of the following glossary terms refer to this illustration:

- DLSE
- DLSE support mode
- Logical DASD subsystem
- Multipath storage director

Bibliography

The manuals listed in the table below contain more detailed information on the subjects discussed in this book. For each manual referenced, the table shows the short and expanded title with the manual's order number, and a short description of its contents.

For information on how to order these manuals, contact your local IBM branch office. Some publications supporting MVS/ESA will be available when their associated product is available.

Short Title	Full Title	Order Number	Contents
Storage Subsystem Library	3380 Manuals		
3380 Direct Access Storage Reference Summary	<i>3380 Direct Access Storage Reference Summary</i>	GX26-1678	Summary card containing 3380 device characteristics
IBM 3380 Direct Access Storage Introduction	<i>IBM 3380 Direct Access Storage Introduction</i>	GC26-4491	Overview of all 3380 models
Maintaining IBM Storage Subsystem Media	<i>Maintaining IBM Storage Subsystem Media</i>	GC26-4495	Description of DASD media maintenance and error handling
Storage Subsystem Library Master Index	<i>Storage Subsystem Library Master Index</i>	GC26-4496	Index to information in 3380 and 3990 manuals
Using the IBM 3380 Direct Access Storage in a VM Environment	<i>Using the IBM 3380 Direct Access Storage in a VM Environment</i>	GC26-4493	Discussion of 3380 use under VM
Using the IBM 3380 Direct Access Storage in a VSE Environment	<i>Using the IBM 3380 Direct Access Storage in a VSE Environment</i>	GC26-4494	Discussion of 3380 use under VSE
Storage Subsystem Library	3990 Storage Control Manuals		
Cache Device Administration	<i>Cache Device Administration</i>	GC35-0101	Describes the AMS cache utility commands necessary to manage cache and to obtain information about cache status and performance
IBM 3990 Storage Control Introduction	<i>IBM 3990 Storage Control Introduction</i>	GA32-0098	Overview of 3990 storage control functions
IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide	<i>IBM 3990 Storage Control Planning, Installation, and Storage Administration Guide</i>	GA32-0100	Detailed information on installation and use of the 3990 storage control
IBM 3990 Storage Control: Reference	<i>IBM 3990 Storage Control: Reference</i>	GA32-0099	Reference information on 3990 functions including descriptions of DASD channel command set
3880 Storage Control Information			
IBM 3880 Storage Control Models 1, 2, 3, and 4 Description	<i>IBM 3880 Storage Control Models 1, 2, 3, and 4 Description Manual</i>	GA26-1661	Overview of 3880 Models 1, 2, 3, and 4 functions
Introduction to IBM 3880 Storage Control Model 13	<i>Introduction to IBM 3880 Storage Control Model 13</i>	GA32-0062	Overview of 3880 Model 13 functions
IBM 3880 Storage Control Model 13 Description	<i>IBM 3880 Storage Control Model 13 Description</i>	GA32-0067	Reference manual for 3880 Model 13 functions
IBM 3880 Storage Control Model 23 Description	<i>IBM 3880 Storage Control Model 23 Description</i>	GA32-0083	Reference manual for 3880 Model 23 functions
IBM 3880 Storage Control Model 23 Installation and Administration Guide	<i>IBM 3880 Storage Control Model 23 Installation and Administration Guide</i>	GA32-0085	Describes how to install and use the 3880 Model 23 effectively
IBM 3880 Storage Control Model 23 Introduction	<i>IBM 3880 Storage Control Model 23 Introduction</i>	GA32-0082	Overview of 3880 Model 23 functions
IBM 3880 Storage Control Model 23 with RPQ #8B0035 Description	<i>IBM 3880 Storage Control Model 23 with RPQ #8B0035 Description</i>	GA32-0087	Reference manual for 3880 Model 23 functions
Introduction to IBM 3880 Storage Control Model 23 with RPQ #8B0035	<i>Introduction to IBM 3880 Storage Control Model 23 with RPQ #8B0035</i>	GA32-0086	Overview of 3880 Model 23 functions
Storage Hardware Maintenance			
EREP User Guide and Reference	<i>Environmental Record Editing and Printing (EREP) User's Guide and Reference</i>	GC28-1378	Description of EREP functions and commands for DASD media reporting

Short Title	Full Title	Order Number	Contents
ICKDSF Primer	<i>Device Support Facilities: Primer for the User of IBM 3380 Direct Access Storage</i>	GC26-4498	Describes how to use ICKDSF with the 3380
ICKDSF User's Guide and Reference	<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033	Description of ICKDSF functions and commands for DASD initialization and maintenance
Physical Planning and Reference Information			
IBM Input/Output Equipment: Installation – Physical Planning for System/360, System/370, and 4300 Processors	<i>IBM Input/Output Equipment: Installation – Physical Planning for System/360, System/370, and 4300 Processors</i>	GC22-7064	Description of physical planning for I/O hardware
IBM Input/Output Equipment: Installation Reference – Physical Planning for System/360, System/370, and 4300 Processors	<i>IBM Input/Output Equipment: Installation Reference – Physical Planning for System/360, System/370, and 4300 Processors</i>	GC22-7069	Description of physical planning for I/O hardware
9370 Information System Installation Manual – Physical Planning	<i>9370 Information System Installation Manual – Physical Planning</i>	GA24-4031	Contains physical planning information for the 9370 family of processors
IOCP User's Guide and Reference	<i>Input/Output Configuration Program User's Guide and Reference</i>	GC28-1027	Shows how to define the I/O configuration data required by the processor complex to control I/O requests, describing the MVS version, the VM version, and the standalone version of IOCP
DATABASE 2 (DB2) Information			
DB2 Release 2 Data Portability	<i>DB2 Release 2 Data Portability</i>	WSCF-8642	Describes moving DB2 data sets and volumes between DASD or systems
IBM DB2 General Information	<i>IBM DATABASE 2 General Information</i>	GC26-4073	Overview of DB2 operation and functions
IBM DB2 Installation	<i>IBM DATABASE 2 Installation</i>	SC26-4084	Describes installation requirements of DB2
IBM DB2 Operations and Recovery Guide	<i>IBM DATABASE 2 Operations and Recovery Guide</i>	SC26-4083	Contains information on designing operating and recovery procedures for DB2.
IBM DB2 Reference	<i>IBM DATABASE 2 Reference</i>	SC26-4078	Describes DB2 utilities and their syntax
IBM DB2 System Planning and Administration Guide	<i>IBM DATABASE 2 System Planning and Administration Guide</i>	SC26-4085	Describes planning activities for installing or migrating to DB2 Release 2 and managing system resources once installed
Data Facility Data Set Services (DFDSS) Information			
DFDSS: General Information (Version 2)	<i>Data Facility Data Set Services: General Information (Version 2)</i>	GC26-4123	Overview of capabilities and requirements
DFDSS: User's Guide and Reference (Version 2)	<i>Data Facility Data Set Services: User's Guide and Reference (Version 2)</i>	SC26-4125	Describes syntax and usage of DFDSS commands
Data Facility SORT (DFSORT) Information			
DFSORT General Information	<i>DFSORT General Information</i>	GC33-4033	Contains introductory material for planners, system support people, managers, or programmers
DFSORT Application Programming Guide	<i>DFSORT Application Programming Guide</i>	GC33-4035	Provides detailed programming information to enable programmers to prepare sort, merge or copy applications
Data Facility Hierarchical Storage Manager (DFHSM) Information			
DFHSM: Installation and Customization Guide	<i>Data Facility Hierarchical Storage Manager: Installation and Customization Guide</i>	SH35-0084	Describes how to install and tailor DFHSM to your needs
DFHSM: System Programmer's Guide	<i>Data Facility Hierarchical Storage Manager: System Programmer's Guide</i>	GH35-0085	Describes the concepts of DFHSM
DFHSM: System Programmer's Reference	<i>Data Facility Hierarchical Storage Manager: System Programmer's Reference</i>	GH35-0083	Describes and explains how to use the DFHSM system programmer, space manager, and operator commands

Short Title	Full Title	Order Number	Contents
Hierarchical Storage Manager (HSM) Information			
HSM Release 3 Installation Guide	<i>Hierarchical Storage Manager Release 3 Installation Guide</i>	GG22-9254	Describes requirements for and activities of HSM installation
OS/VS MVS HSM: General Information	<i>OS/VS MVS Hierarchical Storage Manager: General Information</i>	GH35-0007	Overview of capabilities and requirements
OS/VS2 MVS HSM: User's Guide	<i>OS/VS2 MVS Hierarchical Storage Manager: User's Guide</i>	SH35-0024	Describes syntax and usage of HSM commands
Interactive Storage Management Facility (ISMF) Information			
ISMF User's Guide	<i>MVS/ESA Interactive Storage Management Facility User's Guide</i>	SC26-4508	Describes how to use ISMF to perform data and storage management tasks in an MVS/ESA environment
ISMF User's Guide	<i>MVS/XA Interactive Storage Management Facility User's Guide</i>	G26-4266	Describes how to use ISMF to perform data and storage management tasks in an MVS/XA environment
Data Facility Product (DFP) Information			
DFP: Information	<i>MVS/ESA Data Facility Product Version 3: General Information</i>	GC26-4507	Overview of capabilities and requirements of MVS/DFP
	<i>MVS Extended Architecture Data Facility Product Version 2: General Information</i>	GC26-4142	Overview of capabilities and requirements of MVS/XA DFP
	<i>MVS/370 Data Facility Product: General Information</i>	GC26-4050	Overview of capabilities and requirements of MVS/370 DFP
MVS/DFP Planning Guide	<i>MVS/ESA Data Facility Product Version 3: Planning Guide</i>	SC26-4513	Describes planning for installation of DFP and SMS
	<i>MVS Extended Architecture Data Facility Product Version 2: Planning Guide</i>	GC26-4147	Describes installation of DFP, conversion to integrated catalogs, and conversion to indexed VTOCs in MVS/XA
	<i>MVS/370 Data Facility Product: Planning Guide</i>	GC26-4052	Describes installation of DFP, conversion to integrated catalogs, and conversion to indexed VTOCs in MVS/370
Storage Administration Reference	<i>MVS/ESA Storage Administration Reference</i>	SC26-4514	Explains how to initialize the MVS/DFP Storage Management Subsystem (SMS), maintain it and perform storage management tasks
IMS/VS Information			
Data Base Administration Guide	<i>IMS/VS Version 2 Data Base Administration Guide</i>	SC26-4179	Describes design, implementation and maintenance of IMS data base
General Information Manual	<i>IMS/VS Version 2, General Information Manual</i>	GC26-4180	Overview of IMS concepts
Utilities Reference Manual	<i>IMS/VS Version 2, Utilities Reference Manual</i>	SC26-4173	Provides detailed information on utilities and how to run them
MVS/ESA Planning and Reference Information			
Access Method Services Reference	<i>MVS/ESA Integrated Catalog Administration: Access Method Services Reference</i>	SC26-4500	Describes access method services commands used to manipulate integrated catalog facility catalogs and VSAM data sets
	<i>MVS/ESA VSAM Catalog Administration: Access Method Services Reference</i>	SC26-4501	Describes access method services commands used to manipulate VSAM data sets
Catalog Administration Guide	<i>MVS/ESA Catalog Administration Guide</i>	SC26-4052	Describes how to use the integrated catalog facility
Data Administration Guide	<i>MVS/ESA Data Administration Guide</i>	SC26-4505	Describes how to use access methods (except VSAM) to process data sets
Data Administration: Macro Instruction Reference	<i>MVS/ESA Data Administration: Macro Instruction Reference</i>	SC26-4506	Describes how to code macro instructions for access methods
Initialization and Tuning	<i>MVS/ESA System Programming Library: Initialization and Tuning</i>	GG28-1828	Describes the effects of tuning an MVS/ESA system and how to define these values.

Short Title	Full Title	Order Number	Contents
JCL Reference	<i>MVS/ESA Job Control Language Reference</i>	GC28-1829	Describes the syntax of JCL
JCL User's Guide	<i>MVS/ESA Job Control Language User's Guide</i>	GC28-1830	Describes the use of JCL
MVS Configuration Program Guide and Reference	<i>MVS/ESA MVS Configuration Program Guide and Reference</i>	GC28-1817	Describes use of MVSCP to define the I/O configuration to MVS/ESA
System Generation	<i>MVS/ESA System Generation</i>	GC28-1825	Describes how to do a complete sysgen MVS/ESA
System Messages Volume 1	<i>MVS/ESA Message Library: System Messages Volume 1</i>	GC28-1812	Lists MVS/ESA system messages and responses
System Commands	<i>MVS/ESA Operations: System Commands</i>	GC28-1826	Describes syntax and use of MVS/ESA system commands
MVS/ESA System – Data Administration	<i>MVS/ESA System – Data Administration</i>	SC26-4515	Overview of IBM access methods available in an MVS/ESA system
Service Aids	<i>MVS/ESA System Programming Library: Service Aids</i>	GC28-1844	Describes how to use GTF, LIST, PRDMP, SADMP, and SPZAP
System Messages Volume 2	<i>MVS/ESA System Messages Volume 2</i>	GC28-1813	Lists MVS/ESA system messages and responses
VSAM Administration Guide	<i>MVS/ESA VSAM Administration Guide</i>	SC26-4518	Describes how to create VSAM data sets
MVS/XA Planning and Reference Information			
Catalog Administration Guide	<i>MVS Extended Architecture Catalog Administration Guide</i>	GC26-4046	Describes how to use the integrated catalog facility
Data Administration Guide	<i>MVS Extended Architecture Data Administration Guide</i>	GC26-4140	Describes how to use access methods (except VSAM) to process data sets
Data Administration: Macro Instruction Reference	<i>MVS Extended Architecture Data Administration: Macro Instruction Reference</i>	GC26-4014	Describes how to code macro instructions for access methods
Utilities	<i>MVS/Extended Architecture Data Administration: Utilities</i>	GC26-4150	Describes how to use DFP utility programs for program, device and data management.
Installation: System Generation	<i>MVS/Extended Architecture Installation: System Generation</i>	GC26-4009	Describes how to do a complete sysgen, iogen, or edtgen in MVS/XA
Integrated Catalog Administration: Access Method Services Reference	<i>MVS Extended Architecture Integrated Catalog Administration: Access Method Services Reference</i>	GC26-4135	Describes access method services commands used to manipulate integrated catalog facility catalogs and VSAM data sets
MVS/XA I/O Performance Considerations	<i>MVS/XA I/O Performance Considerations</i>	GG22-9346	Describes the effect of the MVS/XA I/O subsystem on device performance. Tuning considerations are discussed
JCL	<i>MVS/Extended Architecture Job Control Language (JCL)</i>	GC28-1148	Describes syntax and use of JCL in MVS/XA
Message Library: System Messages Volume 1	<i>MVS/Extended Architecture Message Library: System Messages Volume 1</i>	GC28-1376	Lists MVS/XA system messages and responses
Message Library: System Messages Volume 2	<i>MVS/Extended Architecture Message Library: System Messages Volume 2</i>	GC28-1377	Lists MVS/XA system messages and responses
MVS Configuration Program Guide and Reference	<i>MVS/Extended Architecture MVS Configuration Program Guide and Reference</i>	GC28-1335	Describes use of MVSCP to define the I/O configuration to MVS/XA
Operations: System Commands	<i>MVS/Extended Architecture Operations: System Commands</i>	GC28-1206	Describes syntax and use of MVS/XA system commands
System – Data Administration	<i>MVS Extended Architecture System – Data Administration</i>	GC26-4149	Overview of IBM access methods
System Programming Library: Service Aids	<i>MVS/Extended Architecture System Programming Library: Service Aids</i>	GC28-1159	Describes how to use GTF, LIST, PRDMP, SADMP, and SPZAP
VSAM Administration Guide	<i>MVS Extended Architecture VSAM Administration Guide</i>	GC26-4015	Describes how to create VSAM data sets
VSAM Catalog Administration: Access Method Services Reference	<i>MVS Extended Architecture VSAM Catalog Administration: Access Method Services Reference</i>	GC26-4136	Describes access method services commands used to manipulate VSAM data sets

Short Title	Full Title	Order Number	Contents
MVS/370 Planning and Reference Information			
Access Method Services Reference	<i>MVS/370 Integrated Catalog Administration: Access Method Services Reference</i>	GC26-4051	Describes the access method services commands used with VSAM and integrated catalog facility catalogs
	<i>VSAM Catalog Administration: Access Method Services Reference</i>	GC26-4059	Describes the access method services commands used with VSAM
Initialization and Tuning Guide	<i>OS/VS2 MVS System Programming Library: Initialization and Tuning Guide</i>	GC28-1029	Describes how to initialize the system and improve system performance, includes information on GTF
MVS Planning: Global Resource Serialization Guide	<i>OS/VS2 MVS Planning: Global Resource Serialization</i>	GC28-1062	Describes how use the GRS feature in an MVS environment performance, includes information on GTF
Data Administration Guide	<i>MVS/370 Data Administration Guide</i>	GC26-4058	Contains information on using access methods to do input and output
Data Administration Macro Instruction Reference	<i>MVS/370 Data Administration: Macro Instruction Reference</i>	GC26-4057	Describes how to use macros to do input and output
Data Administration: Utilities	<i>MVS/370 Data Administration: Utilities</i>	GC26-4065	Describes how to use IEHLIST to maintain VTOC, IEHMOVE to maintain OS CVOLS, and IEHPROM to protect data sets
JCL User's Guide	<i>MVS/370 JCL User's Guide</i>	GC28-1349	Describes syntax of JCL statements, JES2 and JES3 control statements
System: Data Administration	<i>MVS/370 System: Data Administration</i>	GC26-4056	Describes MVS/370 DFP, and how to modify and extend the data management capabilities of the operating system
System Generation Reference	<i>MVS/370 System Generation Reference</i>	GC26-4063	Describes how to do a sysgen, iogen, or edtgen
VSAM Administration: Macro Instruction Reference	<i>MVS/370 VSAM Administration: Macro Instruction Reference</i>	GC26-4074	Describes using VSAM macro instructions for VSAM data sets
MVS/SP Version 1 General Information Manual	<i>MVS/System Product Version 1 General Information Manual</i>	GC28-1025	Contains overview and planning information for JES3 and JES2 for MVS/370
MVS/SP Version 1 Release 3 Installation Considerations	<i>MVS/System Product Version 1 Release 3 Installation Considerations</i>	GG22-9250	Installation planning considerations for MVS/370
Operator's Library: OS/VS2 MVS System Commands	<i>Operator's Library: OS/VS2 MVS System Commands</i>	GC28-1031	Describes syntax and use of MVS/370 commands
Performance, Availability, and Tuning Information			
An MVS Tuning Perspective	<i>An MVS Tuning Perspective</i>	GG22-9023	Describes the effects of tuning an MVS system
Component Failure Impact Analysis – An Availability Management Technique	<i>Component Failure Impact Analysis – An Availability Management Technique</i>	GC20-1865	Planning for hardware availability through configuration
ISPF/PDF Program Reference	<i>Interactive System Productivity Facility/Program Development Facility Program Reference</i>	SC34-2089	Explains to the system or applications programmer how to use PDF in an MVS environment.
MVS/XA: JES3 User Modifications and Macros	<i>MVS/Extended Architecture: JES3 User Modifications and Macros</i>	SC23-0060	Describes tailoring JES3 exits and macros
MVS/XA RMF Reference and User's Guide	<i>MVS/XA Resource Measurement Facility Reference and User's Guide</i>	LC28-1138	Provides detailed information to operate RMF under MVS/XA
MVS/XA SMF	<i>MVS Extended Architecture System Management Facilities</i>	GC28-1153	Describes how to plan for, install and use SMF to manage the MVS/XA system
OS/VS2 MVS Performance Notebook	<i>OS/VS2 MVS Performance Notebook</i>	GC28-0886	Describes tuning your system to meet performance expectations and optimizing use of your system
OS/VS2 MVS Planning: Global Resource Serialization	<i>OS/VS2 MVS Planning: Global Resource Serialization</i>	GC28-1062	Contains information on how to serialize access to data sets on shared DASD volumes
RACF General Information Manual	<i>Resource Access Control Facility (RACF) General Information Manual</i>	GC28-0722	Provides overview and planning information for the RACF program
OS/VS2 MVS RMF Reference and User's Guide, Version 2	<i>OS/VS2 MVS RMF Reference and User's Guide, Version 2</i>	SC28-0922	Provides detailed information to operate RMF under MVS/370
OS/VS2 System Programming Library System Management Facilities	<i>OS/VS2 System Programming Library System Management Facilities</i>	GC28-1030	Describes how to plan for, install and use SMF to manage the MVS/370 system

Short Title	Full Title	Order Number	Contents
OS/VS2 MVS RMF, Version 2, General Information	<i>OS/VS2 MVS Resource Measurement Facility, Version 2, General Information</i>	GC28-0921	Describes capabilities, functions, and usage of RMF (Version 2 runs on an MVS/370 system)
RMF, Version 3, General Information	<i>Resource Measurement Facility, Version 3, General Information</i>	GC28-1115	Describes capabilities, functions, and usage of RMF (Version 3 runs on an MVS/XA system)
SLR User's Guide	<i>Service Level Reporter User's Guide</i>	SH19-6215	Describes how generate reports using SLR
SLR, Version 2, General Information	<i>Service Level Reporter, Version 2, General Information</i>	GH19-6213	Overview of SLR functions
SPL: Installation, Initialization, and Tuning	<i>System Programming Library: JES2 Installation, Initialization, and Tuning</i>	SC23-0046	Describes requirements for and activities of JES2 installation, initialization, and tuning
	<i>System Programming Library: JES3 Installation, Initialization, and Tuning</i>	SC23-0041	Describes requirements for and activities of JES3 installation, initialization, and tuning
	<i>MVS/ESA System Programming Library: Initialization and Tuning</i>	GG28-1828	Describes the effects of tuning an MVS/ESA system and how to define these values.
Storage Management Library			
MVS SML: Configuring Storage Subsystems	<i>MVS Storage Management Library: Configuring Storage Subsystems</i>	SC26-4409	Describes evaluating hardware configurations, developing capacity plans, and performance, availability and space utilization considerations
MVS SML: Managing Data Sets	<i>MVS Storage Management Library: Managing Data Sets</i>	SC26-4408	Describes managing data sets, catalogs and control data sets. establishing and enforcing data set policy, and data set security
MVS SML: Managing Storage Pools	<i>MVS Storage Management Library: Managing Storage Pools</i>	SC26-4407	Describes storage requirements for groups of data sets, designing storage pools, making transition to pooled storage, and maintaining and monitoring storage pools

Index

- A-BSY-R message 78
- AA4 internal path selection 30
- access method services (IDCAMS) 20
- accounting routines 25
- ACS (Automatic Class Selection)
 - description of 56
- addresses, reserving I/O 5
- administrative planning
 - device delivery, verifying 17
 - project plan 17
 - project team 17
- ALLDATA parameter, DFDSS 93
- ALLEXCP parameter, DFDSS 93
- ALLMULTI, DFDSS parameter 91
- alternate controller selection 31
- alternate path selection 31
- Automatic Class Selection
 - See ACS (Automatic Class Selection)
- Average Block Allocation, using 59
- AVGREC, using 59

- backup and recovery
 - data base 112
 - full-volume 113
 - incremental 113
 - managing 112
 - incremental 113
 - minidisk 112
 - volume 113
 - while migrating data 88
 - restoring data set from tape 88
 - restoring volume from tape 88
- backup, MVS 88
- BDAM data sets, moving with DFDSS 93
- block size
 - selection 58
 - using system determined block size 59
- blocking factor 58
 - for logical record

- cache
 - description 39
- Cache RMF Reporter
 - description 15
 - using 111
- capacity planning
- catalog locking facility
 - available with RACF 52

- catalogs
 - information contained in 20
 - integrated catalog facility 52
 - master 52
 - move guidelines 52
 - OS CVOL 52
 - VSAM 52
- catalogs, moving with DFDSS 89
- channel switches 34
- CHPID statement, MVS IOCP generation 68
- CICS data, moving 93
- CLIST, device dependency in 25
- CNTLUNIT statement 66, 68
 - MVS IOCP generation 3380 Model CJ2
- configuration examples
 - 3380 Model AE4 and AK4 strings attached to 3880 Model 23 42
 - 3380 Model CJ2 string 46
 - 3380 string intermix on 3990 Model 2 or 3 43
 - 3380 4-path strings on 3990 Model 2 or 3 45
 - 4-path and 2-path string intermix on 3990 Model 2 or 3 44
- configuring
 - for availability 37
 - for performance 35
- control units, selection of 39
- copy, full-volume 95
- COPYVOLID parameter, DFDSS 95

- DASD Data Transfer Summary 115
- DASD fast write 40
- DASD, shared 38
- data configuration
 - performance guidelines 54
 - planning 47
- Data Facility Data Set Services 13
- Data Facility Hierarchical Storage Manager 12
- Data Facility SORT 13
- data migration, planning for 47
- data moving
 - backup and recovery during 63
 - documenting backup and recovery while 63
 - moving data to smaller devices 102
 - planning the move sequence 60
 - restoring to unlike device 101
 - scenarios for 60
 - strategy 60
 - tape data sets 102
 - tools
 - using DFHSM 62

- data not supported by SMS 56
- data record
 - optimum blocking for 57
 - selecting a size 58
- data set
 - categories of, MVS 47
 - list of MVS macro and parameter libraries 51
 - list of MVS system data sets 51
 - multiple volume 55
 - performance-critical 54
 - placement considerations 54
 - positioning
 - application data, non-VSAM 53
 - application data, VSAM 53
 - catalogs, integrated catalog facility 52
 - catalogs, master 52
 - catalogs, VSAM 52
 - created by non-IBM programs 54
 - DAM 53
 - DB2 data sets 53
 - ISAM 53
 - JES2 checkpoint data sets 49
 - JES2 spool data sets 48
 - JES3 checkpoint data sets 51
 - JES3 spool data sets 50
 - JES3JCT 51
 - macro and parameter libraries 51
 - multiple volume 53
 - non-VSAM application data 53
 - page and swap 48
 - system data sets, miscellaneous 51
 - tape 54
 - unmovable 53
 - VSAM application data 53
 - work files 48
 - unmovable 55
- data set groups
 - identifying and moving 19
 - performance-critical
 - DB2 data sets 22
 - DFHSM data sets 21
 - identifying 21
 - IMS/VS data sets 21
 - system data sets 21
 - user data sets 22
 - tools for identifying 20
 - access method services 20
 - IEHLIST 20
 - ISMF 20
 - ISPF/PDF 20
 - SLR 20
- data set serialization 38
- data transfer, simultaneous 31
- data volumes, moving with DFDSS 95
- DB2 data sets 22
- DB2 data, moving 93
 - considerations, pre-move 106
 - data move procedure 108
- DB2 data, moving (*continued*)
 - moving data sets or volumes 106
 - storage group considerations 108
- DCB parameters, device-dependent 25
- DELETE parameter, DFDSS 89
- device allocation, obtaining information about 80
- Device End Missing system message 84
- device level selection (DLS)
 - See DLS
- device level selection enhanced (DLSE)
 - See DLSE
- device space considerations 58
- Device Support Facilities
 - See ICKDSF
- device-dependencies
 - identifying and removing 24
- devices, restoring data to unlike 101
- DEVSERV command, in MVS/XA and MVS/ESA 81
- DFDSS 13
 - parameters
 - ALLDATA 93
 - ALLEXCP 93
 - ALLMULTI 91
 - COPYVOLID 95
 - DELETE 89
 - LOGINDDNAME 91
 - LOGINDYNAME 91
 - RECATALOG 89
 - RENAMEUNCONDITIONAL 89
 - TOLERATE(ENQFAILURE) 94
 - using, to move
 - BDAM data sets 93
 - catalogs 89
 - four 3350s to a double capacity 3380 100
 - full volumes 95
 - full-volume copy, example of 96
 - ISAM data sets 93
 - multivolume data sets 91
 - PAM and SAM data sets 93
 - six 3350s to a triple capacity 3380 101
 - system data sets 93
 - two 3350s to one 3380 99
 - unmovable data sets 90
 - VSAM data sets 92
 - 3 standards to one triple capacity 97
- DFHSM
 - commands
 - ADDVOL 62
 - FREEVOL 63
 - MIGRATE 62
 - data sets, frequently accessed 21
 - moving data sets with 62
 - overview 12
- DFSORT 13
- DISPLAY command
 - DISPLAY M, in MVS/XA and MVS/ESA 79
 - DISPLAY M, in MVS/370 79
 - DISPLAY U 78

- DISPLAY command (*continued*)
 - DISPLAY U,,ALLOC 80
 - DLS (device level selection)
 - benefits of 32
 - description of function 32
 - support mode for 3990 32
 - 3380 models supporting 32
 - DLSE (device level selection enhanced)
 - benefits of 33
 - description of function 32
 - support mode for 3990 33
 - 3380 models supporting 32
 - DPS (dynamic path selection)
 - description of functions 31
 - DSCB, copying model 94
 - dual copy 40
 - DUMP/RESTORE, used to move data 101
 - dynamic allocation list 25
 - dynamic path reconnect 32
 - dynamic path selection
 - See DPS
-
- Environmental Record Editing and Printing
 - See EREP
 - EREP
 - description 114
 - reports
 - DASD data transfer summary 115
 - subsystem exception DASD 115
 - system error summary 115
 - error
 - data 115
 - device, permanent 115
 - hardware 115
 - EXCP, device dependencies caused by 25
-
- fast write 40
 - features, prerequisite 6
 - file, backup 112
 - fixes, software 15
-
- General Bill of Form (GBOF) numbers xvi
 - Global Resource Serialization 36
 - GRS 36
 - guest system
 - configuration considerations 38
 - guidelines
 - for space utilization 57
 - hardware capabilities
 - AA4 internal path selection 30
 - alternate path selection 31
 - cache 39
 - channel switches 34
 - data transfer, simultaneous 31
 - DLS 32
 - DLSE 32
 - DPS 31
 - dual copy 40
 - dynamic path reconnect 32
 - fast write 40
 - nonvolatile storage 40
 - speed matching buffer feature 33
 - system-related reserve 31
 - hardware configuration
 - planning 29
 - reviewing
 - RMF measurements 23
 - RMF reports 22
 - understanding 19
 - I/O addresses, reserving 5
 - ICKDSF (Device Support Facilities) 74
 - ICKDSF under MVS 75
 - IEHLIST 20
 - IMS/VS data
 - data move procedure 104
 - general considerations 103
 - IMS/VS data sets 21
 - IMS/VS data, moving 93, 103
 - installation
 - administrative planning 16
 - physical planning
 - cables 5
 - floorspace 5
 - I/O addresses, reserving 5
 - power, electrical 5
 - prerequisite features 6
 - process 65
 - assigning addresses 65
 - defining to MVS 66
 - software planning 7
 - Interactive Storage Management Facility 14
 - internal array and the unit control block 84
 - internal transfer paths
 - Enhanced Subsystem models (AJ4 or AK4 strings) 30
 - Extended Capability models (AD4 or AE4 strings) 30
 - Standard models (AA4 strings) 30
 - introducing the 3380 1
 - IOCP generation, MVS 66
 - IODEVICE statement
 - MVS IOCP generation 68

IODEVICE statement (*continued*)
 3380 Model CJ2 71
 ISAM data sets, moving with DFDSS 93
 ISAM random processing 25
 ISMF 14, 20
 ISPF/PDF 20

JCL, device dependency in 25
 JOBCAT statement 87

like device, definition 13
 logical record
 size and blocking factor 58
 LOGINDDNAME, DFDSS parameter 91
 LOGINDYNAME, DFDSS parameter 91

maintaining media 114
 managing data with SMS 56
 media maintenance 114
 migrating data
 two scenarios 60
 minidisk, backup 112
 missing interrupt, resolving 85
 model DSCB, copying 94
 model groups
 Enhanced Subsystem 1
 Extended Capability 1
 Standard 1
 mount and use attributes 73
 moving IMS/VS data 103
 multipath storage director 33
 multivolume data sets, moving with DFDSS 91

nondisruptive DASD installation 6
 nonvolatile storage 40

O-BSY-R message 78
 OPTCHAN parameter, MVS IOCP generation 68
 optimum blocking for data 57
 OS/VS1 support 11
 out-of-sync condition 84

page data sets, copying 94
 PAM data sets, moving with DFDSS 93
 path group ID, VARY PATH command 84
 paths to a set of devices, obtaining information
 about 79
 paths to devices, obtaining information about 79
 performance
 data placement guidelines 54
 managing 111
 relative seek 33
 skew 36
 tools for monitoring
 channel loading 23
 control unit loading 23
 device loading 23
 VTOC considerations 74
 performance tuning 111
 physical planning 5
 physical record
 selecting a size 58
 planning
 administrative 5
 physical 5
 software 5
 pools, storage 72
 See also volume, pools
 processor
 valid 3380 combinations 2
 program temporary fix (PTF), applying 15
 project plan 17
 project team 17
 publications
 bibliography of 145
 required 16

RECATALOG parameter, DFDSS 89
 record size, optimum 57
 recovery procedures 112
 relative seek performance 33
 reliability
 performance and availability 115
 RENAMEUNCONDITIONAL parameter, DFDSS 89
 reserving a group of paths 31
 Resource Measurement Facility 14
 restoring from backup tape, DFDSS 88
 RMF 14
 measurements 23
 reports, using 22
 RMF reports 22

- SAM data sets, moving with DFDSS 93
- seek performance, relative 33
- selecting physical record size 58
- serialization, data set 38
- Service Information Message 115
- Service Level Reporter 15
- shared DASD 38
- SIM (service information message)
 - alert message, CJ2 115
 - Detail Edit and Summary Reports 117
 - SIM handling procedures 116
 - System Exception Reports 118
- single-path storage director 32
- size of logical record 58
- skew 36
- SLR 15, 20
- smaller devices, moving data to 102
- SMF 15
- SMS (Storage Management Subsystem)
 - data types not supported by 56
 - using to minimize device dependencies 26
- software
 - fixes (PTFs), planning to apply 15
 - planning 5
 - programs required for device support
 - in MVS/ESA 9
 - in MVS/XA 7
 - in MVS/370 10
 - tools
 - Cache RMF Reporter 15
 - DFDSS 13
 - DFHSM 12
 - DFSORT 13
 - ISMF 14
 - release levels, minimum 12
 - RMF 14
 - SLR 15
 - SMF 15
- software changes
 - program temporary fix (PTF)
 - applying 15
- space considerations 58
- space utilization guidelines 57
- speed matching buffer feature (6550) 33
- Start Pending system message 84
- status, device
 - displaying
 - DEVSERV command, in MVS/XA and MVS/ESA 81
 - DISPLAY command, using 77
 - online/offline procedure 82
 - VARY, changing status with 81, 84
 - correcting out-of-sync conditions 84
 - for devices 83
 - for paths 83
 - VARY PATH, using 84
- STEPCAT statement 87
- storage control
 - valid 3380 combinations 2
- storage groups, using 56
- Storage Management Library
 - description 16
- Storage Management Subsystem
 - See SMS (Storage Management Subsystem)
- storage pools 72
 - using
- Storage Subsystem Library
 - library structure xiv
- string 2
 - definition
- strings, configuring 40
- subsystem exception DASD report 115
- swap data sets, copying 94
- switches, channel 34
- system console messages 84
- system data sets 21
- system data sets, moving with DFDSS 93
- System Error Summary 115
- System Exception Reports 118
- system generation, MVS 66
- System Management Facilities 15
- system-managed storage
 - description of 19, 47
- system-related reserve 31
- SYS1.PARMLIB requirements 73
- SYS1.STGINDEX, copying 94

- tape data sets, moving 102
- terminology xii
- TOLERATE(ENQFAILURE), DFDSS parameter 94
- trademark information iii
- transfer paths, internal to 3380
 - See internal transfer paths

- UCB 84
- unit control block 84
- UNIT parameter, MVS system generation 66
- unlike device
 - definition 13
 - restoring data to 101
- unmovable data sets, moving with DFDSS 90
- user data sets 22
- user exit routines 25

- VARY command 81
 - correcting out-of-sync conditions 84
 - for devices 83
 - for paths 83
- VARY command, using the 85
- VARY PATH command, using the 84
- volume
 - backup 112
 - pools 55
 - preparation 73
 - initialization using ICKDSF 74
 - VTOC, size calculation for 73
- volume reserve 31
- volumes, moving with DFDSS 95
- VSAM data sets, moving with DFDSS 92
- VTOC
 - calculating index size 73, 74
 - calculating size 73, 74
 - considerations 74
 - index, performance gain 47
 - information contained in 20
 - size considerations 95

- 2-path string 32
 - description 2
 - illustration 42
- 3380 model groups
 - Enhanced Subsystem 1
 - Extended Capability 1
 - Standard 1
- 4-path string 32
 - description 2
 - illustration 70, 72

Using the IBM 3380 Direct Access Storage
in an MVS Environment

**Reader's
Comment
Form**

GC26-4492-1

This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

If you have applied any technical newsletters (TNLs) to this book, please list them here: _____

Chapter/Section _____

_____ Page No. _____

Comments:

Note: Staples can cause problems with automatic mail-sorting equipment.
Please use pressure-sensitive or other gummed tape to seal this form.

If you want a reply, please complete the following information.

Name _____ Phone No. (_____) _____

Company _____

Address _____

Thank you for your cooperation. No postage is necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail them directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Fold and tape

Please do not staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO. 40 ARMONK, NY



POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Programming Publishing
P.O. Box 49023
San Jose, CA 95161-9023



Fold and tape

Please do not staple

Fold and tape





File Number
S370-07, 4300-07

Printed in U.S.A.

GC26-4492-1

