



**MVS/Extended Architecture
Data Facility Product Version 2:
Planning Guide**

Licensed
Program





**MVS/Extended Architecture
Data Facility Product Version 2:
Planning Guide**

Licensed
Program

| **Third Edition (June 1987)**

| This is a major revision of, and makes obsolete, GC26-4147-1.

| This edition applies to Version 2 Release 3.0 of MVS/Extended Architecture Data Facility Product, Licensed Program 5665-XA2, and to any subsequent releases until otherwise indicated in new editions or technical newsletters.

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 subsequent publication 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 form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, P.O. Box 50020, Programming Publishing, San Jose, California, U.S.A. 95150. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Preface

This publication describes the steps a planner needs to consider before doing the following:

- Installing Data Facility Product
- Installing or converting to integrated catalog facility (ICF) catalogs
- Converting to indexed volume tables of contents (VTOCs)

Organization

This publication contains the following major parts:

- “Part 1. Planning for Data Facility Product Installation” describes the environment required to install Data Facility Product and indicates what information needs to be planned in advance.
 - Chapter 1, “Establishing a Base for Data Facility Product” on page 3, discusses the operating environment required for Data Facility Product, the corequisite and prerequisite program products necessary for installation, and the devices supported on an MVS/Extended Architecture (MVS/XA) Data Facility Product system.
 - Chapter 2, “Determining Storage and Installation Requirements for MVS/XA” on page 17, provides guidelines for estimating storage for Data Facility Product and corequisite program products, and indicates how to plan for the system generation process.
 - Chapter 3, “Interactive Storage Management Facility” on page 21, explains the ISMF environment and describes how to plan for ISMF installation.
 - Chapter 4, “Updating the Publications Library” on page 27, summarizes how the MVS/XA Data Facility Product library differs from OS/VS2 MVS, and describes how to set up the new library.
- “Part 2. Planning for Integrated Catalog Facility Catalogs” explains the integrated catalog facility environment and describes how to plan for catalog installation or conversion.
 - Chapter 5, “Integrated Catalog Facility Catalog Overview” on page 39, describes the integrated catalog facility catalog and its environment.

- Chapter 6, “Integrated Catalog Facility Catalog Installation” on page 45, describes the installation process and the decisions that must be made before installation.
- Chapter 7, “Integrated Catalog Facility Catalog Conversion” on page 63, describes the conversion of OS CVOLs or VSAM catalogs to integrated catalog facility catalogs and indicates how you should prepare for catalog conversion.
- Chapter 8, “Integrated Catalog Facility Catalog Backup and Recovery” on page 77, provides an overview of full catalog recovery and discusses in particular enhanced integrated catalog facility catalog recovery.
- “Part 3. Planning for the Indexed VTOC” describes the structure of the indexed VTOC and outlines how to plan for VTOC conversion.
 - Chapter 9, “Indexed VTOC Overview” on page 89, explains how the indexed VTOC is used.
 - Chapter 10, “Indexed VTOC Conversion” on page 93, describes how to convert VTOCs to indexed VTOCs and indicates how you should prepare for VTOC conversion.
- Appendix A, “Data Facility Product Features” on page 103, lists the features of Data Facility Product that are new or substantially different from previous IBM program products, and indicates where information on these subjects is documented.
- Appendix B, “Using VSAM in an MVS/XA Environment” on page 111, describes the additional capabilities of VSAM for users who run in 31-bit addressing mode, and indicates how to plan for optimal use of 31-bit VSAM support.
- “Glossary of Terms and Abbreviations” on page 113 defines the terms and abbreviations in this book.

An index is also included.

Prerequisite Knowledge

In order to use this book efficiently, you should be familiar with the following topics:

- Software installation procedures
- Data set catalog environments (either OS CVOL or VSAM)
- VTOCs and DASD space utilization
- VSAM (virtual storage access method)

Required Publications

You should be familiar with the information presented in the following publications:

- *MVS/Extended Architecture Installation: System Generation*, GC26-4148, describes the process of system generation and how it relates to installation.
- *MVS/Extended Architecture Catalog Administration Guide*, GC26-4138, describes data set catalogs in greater detail.
- *MVS/Extended Architecture System – Data Administration*, GC26-4149, describes the VTOC in greater detail.
- *MVS/Extended Architecture VSAM Administration Guide*, GC26-4151, describes VSAM and how to use VSAM data sets.

Related Publications

Within the text, references are made to the publications listed in the table below.

Short Title	Publication Title	Order Number
Access Method Services Reference	<i>MVS/Extended Architecture Integrated Catalog Administration: Access Method Services Reference</i>	GC26-4135
	<i>MVS/Extended Architecture VSAM Catalog Administration: Access Method Services Reference</i>	GC26-4136
Assembler H Version 2: General Information	<i>Assembler H Version 2: General Information</i>	GC26-4035
Cache Device Administration	<i>Cache Device Administration</i>	GC35-0101
Catalog Administration Guide	<i>MVS/Extended Architecture Catalog Administration Guide</i>	GC26-4138
Checkpoint/Restart User's Guide	<i>MVS/Extended Architecture Checkpoint/Restart User's Guide</i>	GC26-4139
Conversion Notebook Volume 1	<i>MVS/Extended Architecture Conversion Notebook Volume 1</i>	GC28-1143
Conversion Notebook Volume 2	<i>MVS/Extended Architecture Conversion Notebook Volume 2</i>	GC28-1411

Short Title	Publication Title	Order Number
Cryptographic Unit Support General Information	<i>OS/VS1 and OS/VS2 MVS Cryptographic Unit Support General Information</i>	GC28-1015
CVAF Diagnosis Reference	<i>MVS/Extended Architecture Common VTOC Access Facility Diagnosis Reference</i>	LY26-3958
DADSM Diagnosis Reference	<i>MVS/Extended Architecture DADSM Diagnosis Reference</i>	LY26-3961
DASD Migration Aid General Information	<i>Direct Access Storage Device Migration Aid: General Information</i>	GC26-3972
Data Administration Guide	<i>MVS/Extended Architecture Data Administration Guide</i>	GC26-4140
Data Administration: Macro Instruction Reference	<i>MVS/Extended Architecture Data Administration: Macro Instruction Reference</i>	GC26-4141
Data Facility Product: Customization	<i>MVS/Extended Architecture Data Facility Product Version 2: Customization</i>	GC26-4267
Data Facility Product: General Information	<i>MVS/Extended Architecture Data Facility Product Version 2: General Information</i>	GC26-4142
Data Facility Product: Master Index	<i>MVS/Extended Architecture Data Facility Product Version 2: Master Index</i>	GC26-4146
Device Support Facilities User's Guide and Reference	<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033
DFDSS: General Information	<i>Data Facility Data Set Services: General Information</i>	GC26-4123
DFDSS: User's Guide and Reference	<i>Data Facility Data Set Services: User's Guide and Reference</i>	SC26-4125
DFDSS/ISMF Installation Planning Guide	<i>Data Facility Data Set Services/Interactive Storage Management Facility Installation Planning Guide</i>	SC26-4129
DFHSM System Programmer's Guide	<i>Data Facility Hierarchical Storage Manager: System Programmer's Guide</i>	SH35-0085
DFHSM with ISMF	<i>Data Facility Hierarchical Storage Manager Version 2.2.1 with the Interactive Storage Management Facility</i>	SH35-0108

Short Title	Publication Title	Order Number
EREP User's Guide and Reference	<i>Environmental Recording, Editing, and Printing (EREP) User's Guide and Reference</i>	GC28-1178
Global Resource Serialization	<i>OS/VS2 MVS Planning: Global Resource Serialization</i>	GC28-1062
IBM 3480 Magnetic Tape Subsystem: Introduction	<i>IBM 3480 Magnetic Tape Subsystem: Introduction</i>	GA32-0041
IBM 3480 Magnetic Tape Subsystem: Planning and Migration Guide	<i>IBM 3480 Magnetic Tape Subsystem: Planning and Migration Guide</i>	GC35-0098
IBM 3480 Magnetic Tape Subsystem: User's Guide	<i>IBM 3480 Magnetic Tape Subsystem: User's Guide</i>	GC35-0099
IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility	<i>IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility</i>	SH35-0051
IBM 3800 Printing Subsystem Programmer's Guide	<i>IBM 3800 Printing Subsystem Programmer's Guide</i>	GC26-3846
Initialization and Tuning	<i>MVS/XA System Programming Library: Initialization and Tuning</i>	GC28-1149
IOCP User's Guide and Reference	<i>Input/Output Configuration Program User's Guide and Reference</i>	GC28-1027
ISMF User's Guide	<i>MVS/Extended Architecture Interactive Storage Management Facility User's Guide</i>	GC26-4266
JCL Reference	<i>MVS/Extended Architecture JCL Reference</i>	GC28-1352
JCL User's Guide	<i>MVS/Extended Architecture JCL User's Guide</i>	GC28-1351
Linkage Editor and Loader User's Guide	<i>MVS/Extended Architecture Linkage Editor and Loader User's Guide</i>	GC26-4143
Magnetic Tape Labels and File Structure Administration	<i>MVS/Extended Architecture Magnetic Tape Labels and File Structure Administration</i>	GC26-4145
MVS/SP General Information Manual	<i>MVS/System Product Version 2 General Information</i>	GC28-1118
MVS Configuration Program	<i>MVS/Extended Architecture MVS Configuration Program: Guide and Reference</i>	GC28-1335

Short Title	Publication Title	Order Number
Open/Close/EOV Logic	<i>MVS/Extended Architecture Open/Close/EOV Logic</i>	LY26-3966
RACF General Information Manual	<i>Resource Access Control Facility (RACF): General Information Manual</i>	GC28-0722
RACF Security Administrator's Guide	<i>Resource Access Control Facility Security Administrator's Guide</i>	SC28-1340
SMP System Programmer's Guide	<i>System Modification Program (SMP) System Programmer's Guide</i>	GC28-0673
SMP/E General Information Manual	<i>System Modification Program Extended (SMP/E) General Information Manual</i>	GC28-1106
SMP/E User's Guide	<i>System Modification Program Extended (SMP/E) User's Guide</i>	SC28-1302
Storage Management Reader's Guide	<i>MVS/Extended Architecture Storage Management Library: Storage Management Reader's Guide</i>	GC26-4265
System – Data Administration	<i>MVS/Extended Architecture System – Data Administration</i>	GC26-4149
System Commands	<i>MVS/Extended Architecture Operations: System Commands</i>	GC28-1206
System Generation	<i>MVS/Extended Architecture Installation: System Generation</i>	GC26-4148
System Messages	<i>MVS/Extended Architecture Message Library: System Messages, Volumes 1 and 2</i>	GC28-1376 and GC28-1377
Utilities	<i>MVS/Extended Architecture Data Administration: Utilities</i>	GC26-4150
VSAM Administration Guide	<i>MVS/Extended Architecture VSAM Administration Guide</i>	GC26-4151
VSAM Administration: Macro Instruction Reference	<i>MVS/Extended Architecture VSAM Administration: Macro Instruction Reference</i>	GC26-4152

Summary of Changes

Release 3.0, June 1987

New Program and Data Management Support

Chapter 2, "Determining Storage and Installation Requirements for MVS/XA" on page 17, has been updated to reflect changes in the procedure for defining devices to MVS/XA.

Chapter 3, "Interactive Storage Management Facility" on page 21, has been updated for the volume application of the interactive storage management facility.

Chapter 8, "Integrated Catalog Facility Catalog Backup and Recovery" on page 77, has been added. It includes a description of enhanced integrated catalog facility catalog recovery.

Virtual Storage Constraint Relief

Appendix B, "Using VSAM in an MVS/XA Environment" on page 111, has been updated to reflect changes in VSAM to take advantage of extended architecture and to optimize the use of virtual storage.

Library Updates

Several sections have been added to Chapter 4, "Updating the Publications Library" on page 27 to describe the following:

- Customization restructure of the library
- Diagnosis restructure of the library
- Storage Management Library (SML)
- Storage management family of products
- Storage Subsystem Library (SSL)

Release 2.0, June 1986

Chapter 3, "Interactive Storage Management Facility" on page 21, has been added to support the interactive storage management facility. Information has been added to Chapter 1, "Establishing a Base for Data Facility Product" on page 3, and to Appendix A, "Data Facility Product Features" on page 103, to support the following enhancements to existing program and data management support:

- Interactive storage management facility (ISMF)
- CVAF filter (CVAFFILT)
- Additional block sizes for integrated catalog facility catalogs

Release 1.0, April 1985

Enhancements and New Support

Information to support the following new devices has been added to Chapter 1, "Establishing a Base for Data Facility Product" on page 3, and to Appendix A, "Data Facility Product Features" on page 103:

- IBM 3265 Model 5 Printer
- IBM 3422 Magnetic Tape Subsystem
- IBM 3430 Magnetic Tape Subsystem
- IBM 3480 Magnetic Tape Subsystem
- IBM 3880 Storage Control Model 21
- IBM 3880 Storage Control Model 23
- IBM 3380 Direct Access Storage Models AD4, BD4, AE4, BE4
- IBM 4245 Printer
- IBM 4248 Printer

New programming support includes:

- Catalog address space support
- DADSM support for ERASE on SCRATCH
- Erase on scratch security
- DADSM partial release macro
- DADSM REALLOC enhancements
- Device Console Services
- EXAMINE command
- EXPORT and IMPORT by CIMODE
- IEBCOPY reblocking
- ISO/ANSI/FIPS Version 3 label conversion
- RACF 1.7 tape data set protection
- Tape block count checking
- VSAM JRNAD exits
- VSAM VERIFY support
- 3480 label processing

Version 2 Publications

The Preface includes new order numbers for Version 2.



Contents

Part 1. Planning for Data Facility Product Installation	1
Chapter 1. Establishing a Base for Data Facility Product	3
Operating Systems	3
OS/VS2 MVS	4
MVS/370 (with Data Facility Product 5665-295)	4
MVS/XA (with Data Facility Product 5665-284)	4
Devices	4
Processing Units	5
Input/Output Devices	5
Devices in Compatibility Mode	6
MVS/System Product Job Entry Subsystem (JES2 or JES3)	7
IBM Licensed Programs	9
Requirements for the Installing System	9
Requirements for Data Facility Product and Its Features	11
Compatibility Considerations in a Multisystem Environment	13
Incompatible Programs and Licensed Programs	15
Chapter 2. Determining Storage and Installation Requirements for MVS/XA	17
Data Facility Product	17
Storage Requirements	17
Installation Requirements	17
MVS/System Product JES2 or JES3	18
Corequisite Licensed Programs	18
System Modification Program (SMP/E or SMP)	18
Device Support Facilities	19
Data Facility Data Set Services (DFDSS)	20
Chapter 3. Interactive Storage Management Facility	21
Other Products You Need to Use ISMF	21
Data Sets Used by ISMF	22
Temporary Data Sets (TLIBs)	22
Distribution Data Sets (DLIBs)	22
Execution Data Sets	23
Making ISMF Available to TSO Users	24
ISPF Table Libraries	24
Placement Considerations for the Load Modules	25
Authorization Considerations	26
Chapter 4. Updating the Publications Library	27
MVS/XA Data Facility Product Publications	27
Changes in the MVS/XA Data Facility Product Library	31
Customization Restructure	32

Diagnosis Restructure	32
MVS/XA Storage Management Library	33
Publications for the Storage Management Family of Products	33
Data Facility Data Set Services (DFDSS)	33
Data Facility Hierarchical Storage Manager (DFHSM)	34
Resource Access Control Facility (RACF)	34
Data Facility Sort (DFSORT)	34
Storage Subsystem Library	34
Other Product-Related Publications	35
Prerequisite or Corequisite Licensed Programs	35
Optional Licensed Programs	36

Part 2. Planning for Integrated Catalog Facility Catalogs 37

Chapter 5. Integrated Catalog Facility Catalog Overview	39
The Integrated Catalog Facility Catalog	39
The BCS	40
The VVDS	40
Entries in a Catalog	41
Alternate Master Catalog	42

Chapter 6. Integrated Catalog Facility Catalog Installation	45
An Installation Planning Process	45
The Catalog Environment and the Installation Tasks	46
Conventions for Data Sets and Catalogs	46
Catalog Configuration	48
Catalog Construction	51
Procedures for Data Sets	52
Procedures for Catalogs	54
Security	56
Erasing Residual Data	57
System Resources	58
Virtual Storage Constraint Relief	58
I/O Resources	59
DASD Space and Units	59
Defining the Installation Tasks	60
Scheduling Catalog Conversions	60
The First Production Integrated Catalog Facility Catalog	61
Post-Installation Assessment	61

Chapter 7. Integrated Catalog Facility Catalog Conversion	63
Planning the Conversion of a VSAM User Catalog	63
Checking the CNVTCAT Prerequisites	64
Preparing for CNVTCAT Processing	65
Backing Up VSAM Data Sets, Data Spaces, Catalogs, and OS CVOLs	66
Setting Up the VVDS	66
Setting Up the BCS	67
Converting the Catalog	70
Catalog Conversion with Name Retention	70
Verifying the Conversion	71
Revising Procedures for Integrated Catalog Facility Catalogs	72
Examining Application Requirements	73
Planning the Conversion of a VSAM Master Catalog	73

Planning the Conversion of an OS CVOL	74
Chapter 8. Integrated Catalog Facility Catalog Backup and Recovery	77
Data Error Detection and Diagnosis	77
Full Catalog Recovery	78
Sources of Current Information about Data Sets	78
Full Catalog Recovery Procedure	79
Enhanced Integrated Catalog Facility Catalog Recovery	80
Catalog Locking	81
Locking Authorization	81
Considerations in a Shared Environment	82
Restrictions	82
Recovering a Master Catalog as a User Catalog	83
Recovering Individual Catalog Entries	83
Catalog Backup	84
Steps in the Backup Process	84
Part 3. Planning for the Indexed VTOC	87
Chapter 9. Indexed VTOC Overview	89
Structure	89
Relationship between the VTOC and Its Index	90
VTOC Index Structure	91
Common VTOC Access Facility (CVAF)	91
Chapter 10. Indexed VTOC Conversion	93
Performance Considerations	93
Planning for User Programs That Read VTOCs	93
Determining Which VTOCs Should Be Indexed VTOCs	94
Establishing a VTOC Index at Volume Initialization Time	95
Building a VTOC Index on an Initialized Volume	96
Compatibility	96
Sharing Volumes between Systems	97
Invalidation of a VTOC Index	97
Moving Volumes between Systems	97
Estimating Resource Requirements	98
Mass Storage System Considerations	98
Security	99
Password Protection	99
Resource Access Control Facility (RACF)	100
Authorized Program Facility (APF)	100
DADSM Installation Exits	100
Backup and Recovery	101
Data Set Backup and Recovery	101
Volume Backup and Recovery	101
Appendix A. Data Facility Product Features	103
New Device Support	103
New Programming Support	105
Appendix B. Using VSAM in an MVS/XA Environment	111
Glossary of Terms and Abbreviations	113

Index	119
--------------------	------------

Figures

1. DLIBs Used during SMP ACCEPT Processing	22
2. Execution Data Sets Used during SMP APPLY Processing	23
3. Sample Allocation Parameters for Required Execution Libraries	24
4. Relationship of the BCS and the VVDS	40
5. Integrated Catalog Facility Catalog Relationships	43
6. Catalog Configuration Diagram	49
7. Relationship of a VTOC to Its Index	90

Part 1. Planning for Data Facility Product Installation



Chapter 1. Establishing a Base for Data Facility Product

Chapters 1 through 3 contain information on planning for installation of MVS/Extended Architecture (MVS/XA) Data Facility Product and migration to an MVS/XA environment. An overview of Data Facility Product features and functions is contained in *Data Facility Product Version 2: General Information*. Details on migration to MVS/XA can be found in *Conversion Notebook*.

The planning process assumes that a decision has been made to install Data Facility Product on an existing MVS system. The installation and migration tasks are more complex for non-MVS users; migration to MVS or MVS/XA from OS/VS1, DOS, or other operating systems is not discussed here.

The following terms are used throughout this book:

- OS/VS2 MVS** Refers to the OS/VS2 MVS Release 3.8 operating system with appropriate device and programming support software installed.
- MVS/370** Refers to the OS/VS2 MVS base system with MVS/System Product Version 1, MVS/370 Data Facility Product (5665-295), and appropriate device and programming support software installed.
- MVS/XA** Refers to the OS/VS2 MVS base system with MVS/System Product Version 2, MVS/XA Data Facility Product Version 1 or Version 2 (5665-284 or 5665-XA2), and appropriate device and programming support software installed.
- Sysgen** Refers to the generation of an MVS/XA operating system environment as described in *System Generation*.

For a detailed discussion of the installation process for MVS/XA Data Facility Product, and for the latest information on installation requirements, see the program directory for MVS/XA Data Facility Product.

Operating Systems

This book contains planning information for the installation of MVS/XA Data Facility Product Version 2 Release 3.0. This product must be installed on an OS/VS2 MVS, MVS/370, or MVS/XA base. Any of these three operating systems may be referred to as either the installing system or the generating system.

OS/VS2 MVS

In order to convert an OS/VS2 MVS operating system to MVS/XA, you must first have OS/VS2 MVS (5752-VS2) installed at the Release 3.8 level. Additional programs and licensed programs may also be present on the OS/VS2 MVS Release 3.8 system if you require the functional support they provide (such as support for specific devices).

Products such as Data Facility Device Support (5740-AM7) and Data Facility Extended Function (5740-XYQ), installed on an OS/VS2 MVS system, will be replaced by Data Facility Product. For more information on the functions of these programs, see "Programs Replaced by Data Facility Product" on page 15.

MVS/370 (with Data Facility Product 5665-295)

If you already have installed MVS/370 Data Facility Product (5665-295), you can convert your MVS/370 operating system to MVS/XA by installing MVS/XA Data Facility Product Version 2 Release 3.0 (5665-XA2) and its corequisites. Additional programs and licensed programs may also be present on the MVS/370 system if you require the functional support they provide (such as support for specific devices).

Products such as Data Facility Device Support (5740-AM7) and Data Facility Extended Function (5740-XYQ), installed on an OS/VS2 MVS system, will have already been replaced by MVS/370 Data Facility Product. For more information on the functions of these programs, see "Programs Replaced by Data Facility Product" on page 15.

MVS/XA Data Facility Product (DFP) replaces the functions provided by MVS/370 Data Facility Product. However, before installing MVS/XA DFP, you should delete the MVS/370 DFP modules from your system. Procedures for deleting MVS/370 DFP are described in the program directory for MVS/XA Data Facility Product.

MVS/XA (with Data Facility Product 5665-284)

MVS/XA Data Facility Product Version 2 Release 3.0 may be installed on an MVS/XA Version 1 or Version 2 system. This installation will cause MVS/XA Version 2 Release 3 to replace Version 1 or to replace an earlier release of Version 2.

Devices

A working MVS/XA operating system requires at least one processing unit, printer, console, and system input device, in addition to adequate DASD and tape devices for storage.

Some devices may require a specific level of hardware maintenance in order to function on an MVS/XA system. Contact your IBM representative for

information on hardware maintenance levels and specific devices supported by MVS/XA Data Facility Product.

Processing Units

MVS/XA Data Facility Product can be installed using any processor supported by OS/VS2 MVS, MVS/370, or MVS/XA. This product may also be installed using any processor executing in extended architecture mode.

Execution of the MVS/XA Data Facility Product requires a processor complex, such as the 308x, capable of running in extended architecture mode.

Note that the IBM 3081, 3083, 3084, 3090, and 4381 processor complexes require the Input/Output Configuration Program (IOCP), as distributed with MVS/System Product JES2 or JES3. For more information on IOCP, see *IOCP User's Guide and Reference*.

Input/Output Devices

The table below summarizes the most commonly used IBM I/O devices supported by the MVS/XA operating system. For more information, see *System Generation*.

DASD	Tape	Printer	Console	Other
2305-2	3420	1403-2,-7,-N1	2250-3	2501
3330	3422	3203-5	3251	2540
3333	3430	3211	3277	3505
3340	3480	3262-5	3278	3525
3344		3284	3279	3851
3350		3286		3890
3375		3800-1,-3		
3380		4245		
		4248		

Data Facility Data Set Services (5665-327) can be used to assist in migrating to the IBM 3375 or 3380 Direct Access Storage. For more information, see *DFDSS: General Information*.

Some recent devices require that a specific level of MVS/370 Data Facility Product (5665-295) or Data Facility Device Support (5740-AM7) be present on the generating MVS/370 or OS/VS2 MVS system in order to use these devices during *sysgen*. If you choose to use any of these devices during *sysgen*, your generating system must have the appropriate level of programming support installed. After you have installed MVS/XA Data Facility Product and generated

the MVS/XA system, you may use any of the devices supported under your release of MVS/XA Data Facility Product.

For a list of devices that require Data Facility Device Support if they are to be used during sysgen, see "Programs Replaced by Data Facility Product" on page 15.

Devices in Compatibility Mode

Most devices attached to the system operate in **full function** mode; that is, all features on the device are compatible with and usable on the operating system.

Some devices also operate in **compatibility** mode, which allows you to simulate the function of another device or model. Compatibility mode causes the device to function as a different device, ignoring some or all of the additional features the device might have. This allows you to migrate between devices with minimum impact on user programs.

The devices discussed below are available in either compatibility or full function mode on the MVS/XA system.

3800 Printing Subsystem Model 3

The IBM 3800 Printing Subsystem Model 3 can be used in compatibility mode. It must be generated as a 3800 Model 3 or else it will not have the proper device support, including error recovery support.

The 3800 Model 3 can also be used in full function mode when generated as a 3800 Model 3; full function mode uses all-point addressability and provides access to the advanced printing capabilities of the Model 3. Full function mode for the 3800 Model 3 is sometimes called **page mode** or **all-points-addressable mode**.

To use the full capabilities of the 3800 Model 3, some programming support is required in addition to Data Facility Product. For details on programming requirements, see "Device Support" on page 12. For more information on the functions, see *IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility*.

3480 Magnetic Tape Subsystem

The IBM 3480 Magnetic Tape Subsystem may be generated as a 3480 and used in full function mode under MVS/XA Data Facility Product Version 1 Release 1.2 and Version 2. Under these releases, the 3480 can also be used in 3420 compatibility mode and generated as a **3420C**; this provides only the functions of the 3420, and requires no change to user programs.

The 3420C supports the following 3480 features:

- Tape write-immediate mode (OPTCD = W)
- Full error recovery (ERP) support
- Buffered write mode

The 3420C does **not** support these 3480 features:

- High-speed positioning (block locate)
- Message display feature (MSGDISP macro)
- Dynamic device reconfiguration (DDR)

The 3480 in full function mode, and the 3480 in 3420 compatibility mode, cannot both be generated in the same sysgen on the same MVS/XA system. If you are generating a 3480 in full function mode, you cannot maintain compatibility mode for any other 3480s on the system. However, catalog entries for both the 3480 and the 3420C are synonymous, and both entries may exist in the same catalog. This allows you to share catalogs among systems or migrate to a full function 3480 without changing catalog entries. Each system will use the catalog entry appropriate to the 3480 device type that was generated on that system.

If you plan to use the 3480 in 3420 compatibility mode, you should define an esoteric name for the device (using the UNITNAME sysgen macro) and use this name in the UNIT parameter of the DD cards in your JCL. This will allow you to install the full function 3480 without changes to JCL.

For more information on the functions of the 3480 in full function or compatibility mode, see *IBM 3480 Magnetic Tape Subsystem: Introduction*.

MVS/System Product Job Entry Subsystem (JES2 or JES3)

In order to sysgen MVS/XA Data Facility Product, you must have MVS/System Product (MVS/SP) Version 2 distribution libraries installed. MVS/SP provides essential system support and program distribution libraries for the MVS/XA operating system environment.

MVS/XA Data Facility Product Version 2 Release 3.0 requires Version 2 Release 2.0 of MVS/SP JES2 (5740-XC6), or MVS/SP JES3 (5665-291).

The tables below list which releases of MVS/SP JES2 or JES3 are required to support specific devices or Data Facility Product features. Any later release of MVS/SP Version 2 will also contain the appropriate support.

A specific release of Data Facility Product and/or other licensed programs may be required in addition to MVS/SP for complete support. For more information about software requirements for Data Facility Product features, see "Requirements for Data Facility Product and Its Features" on page 11.

Device	MVS/SP JES2 or JES3 Version
3262 Printer Model 5	JES2 2.1.2 or JES3 2.1.2
3422 Magnetic Tape Subsystem	JES2 2.1.2 or JES3 2.1.2
3430 Magnetic Tape Subsystem	JES2 2.1.2 or JES3 2.1.2
3480 Magnetic Tape Subsystem (full function)	JES2 2.1.2 or JES3 2.1.2
3480 Magnetic Tape Subsystem (compatibility)	JES2 2.1.2 or JES3 2.1.2
3800 Printing Subsystem Model 3 (full function)	JES2 2.1.2 or JES3 2.1.2
3800 Printing Subsystem Model 3 (Model 1 compatibility)	JES2 2.1.2 or JES3 2.1.2
3380 Direct Access Storage Models AD4 and BD4	JES2 2.1.1 or JES3 2.1.1
3380 Direct Access Storage Models AE4 and BE4	JES2 2.1.2 or JES3 2.1.2
3880 Storage Control Model 11	JES2 2.1.0 or JES3 2.1.0
3880 Storage Control Model 13	JES2 2.1.0 or JES3 2.1.0
3880 Storage Control Model 21	JES2 2.1.2 or JES3 2.1.2
3880 Storage Control Model 23	JES2 2.1.2 or JES3 2.1.2
4245 Printer	JES2 2.1.2 or JES3 2.1.2
4248 Printer	JES2 2.1.2 or JES3 2.1.2

DFP Feature	MVS/SP JES2 or JES3 Version
ISO/ANSI/FIPS tape labels and file structure	JES2 2.1.2 or JES3 2.1.2 with the JES3 ANSI support feature (FMID JJS2351)
RACF generic profile support	JES2 2.1.2 with PTF UZ90212 (or UZ90285 for users of RACF Release 6); or JES3 2.1.2 with PTF UZ90212 (or UZ90285 for users of RACF Release 6)

For a discussion of the difference between JES2 and JES3, and the differences among the various releases of MVS/System Product, consult *MVS/SP Version 2 General Information Manual*.

IBM Licensed Programs

Requirements for the Installing System

Before you can install MVS/XA Data Facility Product on an OS/VS2 MVS or MVS/370 base, you must have the following products (or their equivalents) installed as part of the generating system.

Assembler H Version 2

To perform the assemblies needed for system generation and maintenance of MVS/XA Data Facility Product, Assembler H Version 2 (5668-962) is required. This version of Assembler H contains support for 31-bit addressing in assembler language programs.

For more information on Assembler H Version 2, see *Assembler H Version 2: General Information*.

System Modification Program (SMP/E or SMP)

In order to install Data Facility Product, System Modification Program Extended (SMP/E) (5668-949) or System Modification Program (SMP) Release 4 is required. For more information on use of the System Modification Program (SMP/E or SMP), see "System Modification Program (SMP/E or SMP)" on page 18.

Environmental Recording, Editing, and Printing (EREP)

At least Version 3 Release 2 of the Environmental Recording, Editing, and Printing (EREP) licensed program (5658-260) is required for error analysis and recording under MVS/XA.

For more information on EREP, see *EREP User's Guide and Reference*.

Device Support Facilities (ICKDSF)

To initialize DASD volumes, to write IPL text, and to create indexed VTOCs for your system, you must use the MVS/XA version of Device Support Facilities (5655-257) or its equivalent. Support for IBM 3380 Direct Access Storage Models AD4, BD4, AE4, and BE4 requires the functions provided by Release 7 with PTF UZ90326, or ICKDSF Release 8. To support the IBM 3380 Direct Access Storage attached to the IBM 3880 Storage Control Model 13, Release 6 of Device Support Facilities is required. For all other DASDs, any level of Device Support Facilities that is compatible with your DASD may be used on the generating system.

Only Device Support Facilities Release 6 and later levels should be used in the MVS/XA Data Facility Product environment. Previous levels of Device Support Facilities will be removed during the installation of MVS/XA Data Facility Product.

For more information on Device Support Facilities, see *Device Support Facilities User's Guide and Reference*.

Data Facility Data Set Services (DFDSS)

To back up DASD files on tape, in a format that can be restored after installation of MVS/XA Data Facility Product, Data Facility Data Set Services or its equivalent is required.

Because any other IBM program that may have been used to back up DASD files on tape will be deleted when Data Facility Product is installed, you will not be able to use these programs to restore the tapes they created. Furthermore, you will not be able to use DFDSS to restore those tapes, because DFDSS can restore only data sets that it has dumped. You should therefore use DFDSS or its equivalent to back up your OS/VS2 MVS distribution libraries and all other data sets you want to save **before** installing Data Facility Product, so that these may be restored if necessary on the new MVS/XA system. For a discussion of alternatives to previous IBM dump and restore programs, see "Programs Not Supported by Data Facility Product" on page 16.

DFDSS Version 2 Release 3 provides support for doing the following through the interactive storage management facility (ISMF) of Data Facility Product:

- Copying data, either by data set or by full volume, to another volume
- Dumping data, either by data set or by full volume, to another medium
- Restoring data, logically or physically, to a volume
- Releasing allocated, but unused, space from all data sets on a volume
- Compressing unused space from partitioned data sets on a volume
- Reducing fragmentation on a volume

DFDSS Version 2 Release 3 also supports expiration dates greater than 1999, linear data sets, VSAM catalog aliases, and recovery of damaged integrated catalog facility catalogs.

DFDSS Version 2 Release 2 or later is required to support logical data set restore and data set copy of VSAM data sets and catalogs with increased VSAM block sizes. DFDSS Version 2 Release 1 and Version 1 Release 2.1 require toleration PTFs for the increased block sizes; these PTFs return a reason code stating that a data set with a newly allowed block size must be restored in the environment in which it was created. See Data Facility Product *Program Directory* for more information about these toleration PTFs.

For more information on DFDSS, see *DFDSS: General Information*.

Interactive System Productivity Facility (ISPF)

To use the interactive storage management facility (ISMF) of Data Facility Product, Interactive System Productivity Facility Version 2 Release 2 (5665-319) is required. The edit and browse features of ISMF require Version 2 Release 2 (5665-317) of Interactive System Productivity Facility/Program Development Facility (ISPF/PDF).

Time Sharing Option Extensions (TSO/E)

TSO/E is the base for all TSO enhancements and provides improved interactive time sharing. The interactive storage management facility (ISMF) of Data Facility Product requires Version 1 Release 2 (5665-285) of Time Sharing Option Extensions (TSO/E). To use ISPF tables at addresses over 16 megabytes, TSO/E Version 1 Release 2.1 (5665-285) is required.

Requirements for Data Facility Product and Its Features

Data Facility Product Base

The licensed programs listed in the table below are optional in the MVS/XA environment. However, if installed with Data Facility Product, these programs must meet the minimum release level indicated in the table in order to ensure product compatibility.

Licensed Program	Number	Release
Advanced Communications Function for TCAM (ACF/TCAM) Version 2	5735-RC3	4.0
Advanced Communications Function for VTAM (ACF/VTAM) Version 2 or 3	5665-280 or 5665-289	1.0
BTAM/System Product (BTAM/SP)	5665-279	1.0
Customer Information Control System (CICS/VS) Version 1	5740-XX1	6.0
Data Facility Hierarchical Storage Manager (DFHSM) Version 2	5665-329	2.1
Data Facility Sort (DFSORT)	5740-SM1	5.0
Information Management System/Virtual Storage (IMS/VS) Version 1	5740-XX2	2.0 plus compatibility feature
Resource Access Control Facility (RACF) Version 1	5740-XXH	7.0
Resource Measurement Facility (RMF) Version 3	5665-274	1.0

Note: Different levels of these products may also be required to support specific features and devices under MVS/XA Data Facility Product. For a list of features requiring additional software support, see "Data Facility Product Features" on page 12. For a list of devices requiring additional software support, see "Device Support" on page 12.

Chapter 4, "Updating the Publications Library" on page 27, contains a list of publications related to these licensed programs.

Data Facility Product Features

The licensed programs listed in the table below are required if Data Facility Product is to support the listed feature. For some features, a specific level of MVS/SP JES2 or JES3 is also required. For more information, see "MVS/System Product Job Entry Subsystem (JES2 or JES3)" on page 7.

Feature	Products Required	Program Number
Accessing DFDSS copy, dump, restore, compress, and release functions for a data set through ISMF	Data Facility Data Set Services (DFDSS) Version 2 Release 2	5665-327
Accessing DFDSS copy, dump, restore, compress, release, and defrag functions for volumes through ISMF	Data Facility Data Set Services (DFDSS) Version 2 Release 3	5665-327
Accessing DFHSM functions through ISMF	Data Facility Hierarchical Storage Manager (DFHSM) Version 2 Release 2.1	5665-329
Edit and browse features of ISMF	Interactive System Productivity Facility/Program Development Facility (ISPF/PDF) Version 2 Release 2	5665-317
Enciphering and deciphering data using access method services	Programmed Cryptographic Facility or Cryptographic Unit Support (<i>See Note</i>)	5740-XY5 or 5740-XY6
Using ISPF tables at addresses over 16 Mb	Time Sharing Option Extensions (TSO/E) Version 1 Release 2.1	5665-285

Note: If you include the Cryptographic Unit Support licensed program in your system, you must also include its hardware prerequisites. For more information, see *Cryptographic Unit Support General Information*.

Different levels of these products may also be required to support specific devices under MVS/XA Data Facility Product. For a list of devices that require additional software support, see "Device Support."

Chapter 4, "Updating the Publications Library" on page 27, contains a list of publications related to these licensed programs.

Device Support

The following table lists the additional licensed programs that are required for complete device support. For some devices, a specific level of MVS/SP JES2 or JES3 is also required. For more information, see "MVS/System Product Job Entry Subsystem (JES2 or JES3)" on page 7.

Device	Products Required	Program Number
3380 Direct Access Storage Models AD4, BD4, AE4, BE4	ICKDSF Release 7 with PTF UZ90326, or Release 8	5655-257
3480 Magnetic Tape Subsystem (full function)	DFHSM Version 1 Release 3.1 (with PTF), or Version 2 Release 1.0	5740-XRB 5665-329
	DFSORT Release 5 (with PTF)	5740-SM1
3800 Printing Subsystem Model 3 (full function)	Print Services Facility Release 1.0	5665-275

Note: Wherever possible, program temporary fix (PTF) numbers are included in the table when a PTF tape is required for device or programming support. If the letters "PTF" are indicated without a number, the PTF number was not available at the time this manual was published. For detailed information on PTF numbers, see the Data Facility Product *Program Directory*.

Compatibility Considerations in a Multisystem Environment

In a multisystem environment, prior releases of Data Facility Product cannot function with certain features of later releases unless toleration PTFs have been installed.

In particular, toleration PTFs are needed on the following systems to make them compatible with the following features:

Feature	MVS/XA DFP 2.2.0	MVS/XA DFP 2.1.0	MVS/XA DFP 1.1.2	MVS/370 DFP 1.1.1
Increased block sizes	No PTF	PTF	PTF	PTF
Expiration dates beyond 1999	PTF	PTF	PTF	PTF
Export of aliases	PTF	PTF	PTF	PTF
Linear data sets	PTF	PTF	PTF	PTF

Note: You must consult the install section of the current Preventive Service Planning (PSP) Facility for MVS/XA Data Facility Product Version 2 Release 3.0 to obtain the toleration PTF numbers, requisite PTF numbers, and status information.

To retrieve the information from the PSP Facility for this release of MVS/XA Data Facility Product Version 2, contact the IBM Support Center and specify the following information:

Release	Upgrade	Subset
MVS/XA DFP 2.3.0	MVSXADFP230	HDP2230

Increased Block Sizes

MVS/XA Data Facility Product Version 2 Release 2 has expanded support for block sizes of 512, 1K, 2K, and 4K bytes to include multiples of 512 bytes for blocks 512 through 8K bytes, and multiples of 2K bytes for blocks 8K through 32K bytes in length. The toleration PTFs for this feature allow users to process data sets created on MVS/XA Data Facility Product Version 2 Release 2, but not to create them.

In addition, when data sets have been (1) created in MVS/370 Data Facility Product Version 1 Release 1.1, MVS/XA Data Facility Product Version 1 Release 1.2, or MVS/XA Data Facility Product Version 2 Release 1, (2) redefined in MVS/XA Data Facility Product Version 2 Release 2, and then (3) used by the earlier releases, more virtual storage may be needed because new block sizes may have been assigned by that redefinition.

| Expiration Dates beyond 1999

| With the installation of MVS/System Product Version 2 Release 2 and MVS/XA Data Facility Product Version 2 Release 3.0, expiration dates may be specified that extend up to and include the year 2155. To be able to move or share catalogs and data sets between Version 2 Release 3.0 and the prior releases of Data Facility Product, you will need to install the appropriate PTFs for expiration dates beyond 1999.

| Export of Aliases

| As of MVS/XA Data Facility Product Version 2 Release 3.0, when an integrated catalog facility catalog is exported, any aliases defined for the catalog are exported with the catalog. To be able to import these catalogs on the prior Data Facility Product releases, you will need to install the appropriate PTFs for export of aliases.

| Linear Data Sets

| MVS/XA Data Facility Product Version 2 Release 3.0 introduces, in addition to entry sequence, key sequence, and relative record, a fourth kind of VSAM data organization: linear. A linear data set is a continuous string of data with no control information.

| To be able to process linear data sets on the prior Data Facility Product releases specified above, you will need to install the appropriate PTFs for linear data sets. These PTFs allow the prior releases to process a linear data set in control interval mode and to receive return and reason codes for logical errors that occur while a linear data set is being processed.

Incompatible Programs and Licensed Programs

Programs Replaced by Data Facility Product

The functions of the licensed programs listed below are incorporated in Data Facility Product:

- Data Facility Device Support (DFDS) (5740-AM7)
- Data Facility Extended Function (DFEF) (5740-XYQ)
- Sequential Access Method-Extended (SAM-E) (5740-AM3)
- Access Method Services Cryptographic Option (5740-AM8)
- Offline IBM 3800 Utility (5748-UT2)

During installation, Data Facility Product will delete these licensed programs (if they are installed on your system) and replace their respective modules. After Data Facility Product has been installed, you may not install or reinstall any of the licensed programs listed above.

The table below indicates which levels of the replaced licensed programs contain support for various functions. MVS/XA Data Facility Product Version 2 Release 3.0 supports these functions.

Program	Release	Functions Supported
Data Facility Device Support	1.0	Indexed VTOC Alternate master catalog 3375, 3380 Direct Access Storage
	1.5	3880 Storage Control Model 11 3880 Storage Control Model 13
	1.6	3800 Printing Subsystem Model 3 in Model 1 compatibility mode
Data Facility Extended Function	1.0	Integrated catalog facility catalogs VSAM extensions
SAM-E	1.0	SAM extensions
Access Method Services Cryptographic Option	—	Enciphering and deciphering of data
Offline 3800 Utility	—	3800 Printing Subsystem Model 1 operating in offline mode

Programs Not Supported by Data Facility Product

In addition to those listed in the previous section, there are several programs that are not supported by Data Facility Product and that will not function in the MVS/XA environment. These programs, and their suggested alternatives, are listed in the table below.

Old Program	Function	Alternative
Analysis Program-1 (AP-1)	Analyze DASD errors	Device Support Facilities
Direct Access Storage Dump Restore (DRWDASDR)	A licensed program used to dump and restore disks	Data Facility Data Set Services (DFDSS) Note that neither dump format produced by DRWDASDR can be restored by DFDSS.
IBCDASDI utility	Initialize DASD	Device Support Facilities
IBCDMPRS utility	Dump and restore disks in a stand-alone environment	Data Facility Data Set Services (DFDSS): restore functions only
IEHDASDR utility	Initialize disks and dump or restore disks in a VS1 or VS2 MVS environment	Device Support Facilities: disk initialization Data Facility Data Set Services (DFDSS): dump/restore

Chapter 2. Determining Storage and Installation Requirements for MVS/XA

The MVS/XA operating system requires storage for Data Facility Product, for MVS/System Product JES2 or JES3, for all licensed programs installed on the system, and for the various distribution and target libraries supported by each product. Storage requirements will vary depending upon the products installed on your system and the DASD you choose to store on. Consult the General Information and/or Installation manuals for your products, and the program directories for Data Facility Product and MVS/SP for detailed information on storage requirements.

The procedure for generating an MVS/XA system may vary depending upon the release level of Data Facility Product you are installing, and the types of devices you plan to generate on your system. A thorough discussion of system generation requirements can be found in the program directory for Data Facility Product. More details on the system generation process will be found in *System Generation*.

Data Facility Product

Storage Requirements

Detailed information about storage requirements for MVS/XA Data Facility Product Version 2 will be found in the *Program Directory*.

Installation Requirements

Assembler H Version 2, SMP/E or SMP Release 4, and the MVS/XA linkage editor (as provided with MVS/XA or MVS/370 Data Facility Product) are required to generate the MVS/XA system.

An MVS/XA sysgen can be performed on a system with MVS/SP JES2 1.3.0 or JES3 1.3.1 installed in order to migrate from MVS/370 to MVS/XA. Testing of that MVS/XA system requires a processor IMLed in extended architecture mode. The VM/XA Migration Aid licensed program (5664-169) can also be used to assist in migrating to MVS/XA.

The MVS/XA installation process is summarized in the following table.

Installing MVS/XA DFP Using SMP 4	Installing MVS/XA DFP Using SMP/E
<ol style="list-style-type: none"> 1. Update your existing distribution libraries for Data Facility Product, using the ACCEPT function of SMP. 2. Execute a Stage I sysgen with GENTYPE = ALL. 3. Execute Stage II of a complete sysgen. New devices may be added to the system at this time. 	<ol style="list-style-type: none"> 1. Update your existing distribution libraries for Data Facility Product, using the ACCEPT function of SMP/E. 2. Execute a Stage I sysgen with GENTYPE = ALL. 3. You may skip Stage II by using the JCLIN and APPLY functions of SMP/E to update your system libraries. To install any new devices, execute the MVS configuration program (MVSCP) to update your configuration. <p>If you choose not to skip Stage II, new devices may be added to the system at this time.</p>

User routines, macros, and programs that do not have sysgen support will have to be reinstalled after installation of Data Facility Product. This will also be true for non-IBM licensed programs that are not supported by the system generation process.

For details on planning and performing a complete sysgen, see *System Generation*. For information about how to use the MVS configuration program, see *MVS Configuration Program*.

MVS/System Product JES2 or JES3

Information on storage and sysgen requirements for MVS/SP JES2 or JES3 is contained in the program directory for MVS/SP JES2 or MVS/SP JES3.

Corequisite Licensed Programs

System Modification Program (SMP/E or SMP)

The System Modification Program (either SMP/E or SMP) is required to incorporate new products and service into your system libraries. Either of the following may be used:

- System Modification Program Extended (SMP/E)
- System Modification Program (SMP) Release 4

SMP/E contains all the functions of SMP, with the addition of several new functions.

Either SMP/E or SMP must be installed on the generating system before the MVS/XA system can be installed. In addition, modifications to the new MVS system should be made in the SMP/E or SMP installation format, in order to simplify future maintenance.

For more information on SMP/E and SMP, see *SMP/E General Information Manual*, or *SMP System Programmer's Guide*.

Storage is required for SMP/E or SMP Release 4, as well as for the SMP/E or SMP data sets required for installation of Data Facility Product. Storage estimates for the SMP/E or SMP product are provided in the *SMP/E or SMP Program Directory*.

Storage requirements for SMP data sets used in the Data Facility Product installation process are discussed in the *Program Directory* for Data Facility Product.

Note that installation of Data Facility Product requires a large amount of space in the SMP data sets. To ensure that sufficient space is available, you should plan to perform a space analysis of your SMP data sets prior to installation and between installation steps.

Because of the large number of macros distributed with Data Facility Product, you may want to use a temporary SMPMTS data set during Data Facility Product installation; it can then be deleted after installation is complete. For more information on using a temporary SMPMTS data set, see the Data Facility Product *Program Directory*.

Device Support Facilities

The Device Support Facilities system control program requires approximately 512K bytes of storage in which to operate. See the Device Support Facilities *Program Directory* for details on storage requirements for Device Support Facilities.

Before sysgen, you must initialize the DASD volumes that are to contain the new system data sets. DASD volume initialization, as performed by Device Support Facilities, includes the following:

- Varying the volumes offline
- Executing the Device Support Facilities commands INIT, BUILDIX, and REFORMAT to write home addresses, volume labels, VTOCs (with or without an index), and IPL text on the volumes
- Mounting the volumes

Device Support Facilities User's Guide and Reference describes how to initialize DASD volumes under OS/VS2 MVS and MVS/XA.

Data Facility Data Set Services (DFDSS)

The Data Facility Data Set Services licensed program requires a minimum of 430K bytes of virtual storage in which to operate. The exact amount of virtual storage required depends upon the operation performed with DFDSS, and ranges from approximately 430K to 900K bytes. For details on how to estimate the amount of storage a specific DFDSS operation will require, see "Storage Requirements—System DFDSS" in *DFDSS: User's Guide and Reference*.

Chapter 3. Interactive Storage Management Facility

The interactive storage management facility (ISMF) helps you manage data and storage interactively. It is designed to use the space management and backup/recovery functions provided by Data Facility Hierarchical Storage Manager (DFHSM) and Data Facility Data Set Services (DFDSS) to perform a variety of tasks.

As an ISPF application, ISMF has a structure that is modeled after ISPF. ISMF is menu-driven with fast paths for many of its functions. It can be used successfully by both experienced and inexperienced users. ISMF can be customized to fit the needs of the individual installation, or of the users within an organization. Like ISPF, ISMF supports split screen mode. With split screen you can run two independent dialogs at one time (ISPF, ISMF, or both). If you use IBM 3290 terminals, you can run up to four dialogs at a time. ISMF uses the ISPF log to record errors and successful completion of tasks. ISMF functions can be restricted using the program control feature of RACF 1.7. For more information on tailoring ISMF to your needs, refer to *Data Facility Product: Customization* and *ISMF User's Guide*.

Other Products You Need to Use ISMF

To use ISMF, you must have the following products installed on your system, at the given release level or higher:

- Interactive System Productivity Facility (ISPF) Version 2 Release 2
- Time Sharing Option Extensions (TSO/E) Version 1 Release 2

To make full use of ISMF's storage management and authorization functions, and the ISPF edit and browse functions, you need:

- Interactive System Productivity Facility/Program Development Facility (ISPF/PDF) Version 2 Release 2
- Time Sharing Option Extensions (TSO/E) Version 1 Release 2.1
- Data Facility Hierarchical Storage Manager (DFHSM) Version 2 Release 2.1
- Data Facility Data Set Services (DFDSS) Version 2 Release 3
- Resource Access Control Facility (RACF) Version 1 Release 7

Data Sets Used by ISMF

Before you can begin to plan for ISMF, you need to understand how to use the following data sets:

- Temporary data sets (TLIBs)
- Distribution data sets (DLIBs)
- Execution data sets

For additional information on planning for installation and execution of ISMF, see *DFDSS/ISMF Installation Planning Guide* and *DFHSM with ISMF*.

Temporary Data Sets (TLIBs)

ISMF is received on tape in SMP RELFILE format. SMP RECEIVE processing takes the RELFILEs off the tape and puts them in temporary data sets (TLIBs). SMP RECEIVE processing automatically allocates the space for the TLIBs and names the data sets for you. If you choose, you can increase the size of the data set or modify the data set name by changing the prefix. For more information on changing the prefix or allocating parameters, refer to the *Program Directory*. You will indicate where to put the TLIBs when you specify volume information.

Distribution Data Sets (DLIBs)

Distribution data sets (DLIBs) are used for maintenance of ISMF. Figure 1 lists the DLIBs that are used during SMP ACCEPT processing. You can defer the allocation of these data sets until that time.

Distribution Data Set Name	Distribution Library	DDNAME
SYS1.ADGTLLIB	ISMF load module library	ADGTLLIB
SYS1.ADGTPLIB	ISMF panel library	ADGTPLIB
SYS1.ADGTMLIB	ISMF message library	ADGTMLIB
SYS1.ADGTTLIB	ISMF table library	ADGTTLIB
SYS1.AGENLIB	MVS sysgen library	AGENLIB

Figure 1. DLIBs Used during SMP ACCEPT Processing

SMP ACCEPT processing loads the temporary data sets to the following distribution libraries:

- Load module library contains the executable load modules for ISMF.
- Panel library contains data entry panels, online help panels, and confirmation panels.
- Message library contains the text for short and long messages.
- Table library contains the ISMF command tables and the application profile pool tables.
- Sysgen library contains the sysgen macros for ISMF.

For more information, refer to *System Generation, DFSS/ISMF Installation Planning Guide*, and *DFHSM with ISMF*.

The *Program Directory* distributed with DFP 2.3.0 contains the most accurate and current information on requirements for the various libraries. You may want to allocate more than the required amount of space to allow for future expansion.

Where to Put the DLIBs

You need to consider some important issues before deciding where to put the distribution libraries. These include cost, performance, and recovery.

Cost: The DLIBs are used for maintenance purposes only. Therefore, you should put them on a direct access storage device that is used for long-term storage. The main concern with storing the DLIBs is to ensure that they can be recovered easily.

Performance: Performance is not a major concern for DLIBs. Rather, it is important that they have long-term integrity because they are used only for maintenance and backup.

Recovery: It is recommended that you put the new DLIBs with your other DLIBs so you can apply your existing backup and recovery procedures.

Execution Data Sets

Execution data sets, or target libraries, are required to execute ISMF. SMP APPLY processing uses the data sets listed in Figure 2.

Execution Data Set Name	Execution Library	DDNAME
SYS1.DGTLLIB	ISMF load module library	DGTLLIB
SYS1.DGTPLIB	ISMF panel library	DGTPLIB
SYS1.DGTMLIB	ISMF message library	DGTMLIB
SYS1.DGTTLIB	ISMF table library	DGTTLIB

Figure 2. Execution Data Sets Used during SMP APPLY Processing

How Much Space to Allocate for Execution Libraries

Figure 3 on page 24 shows the sample allocation parameters for the execution libraries that you need to allocate. This includes record format, logical record length, block size, total space required (shown both as number of blocks and equivalent number of IBM 3380 tracks), and the number of partitioned data set directory blocks required.

Note: The figures provided are for planning purposes only. The *Program Directory* distributed with the product contains the most accurate and current information on requirements for the various libraries.

Execution Data Set Name	RECFM	LRECL	BLKSIZE	Number of Blocks	3380 Tracks	PDS Directory Blocks
SYS1.DGTLLIB	U	—	6144	110	12	6
SYS1.DGTPLIB	FB	80	3120	180	10	15
SYS1.DGTMLIB	FB	80	3120	80	7	10
SYS1.DGTTLIB	FB	80	3120	4	1	1

Figure 3. Sample Allocation Parameters for Required Execution Libraries

Where to Put the Execution Libraries

You need to consider some important issues before deciding where to put the execution libraries. These include cost, performance, and recovery.

Cost: The execution data sets are used for execution purposes only. Therefore, you may want to put them on a direct access storage device that is used for high performance and quick recovery. However, performance should be weighed against the amount of usage you expect for the execution data sets.

Performance: Performance is the primary concern in placing execution data sets. If you choose to make the load modules part of the link pack area, you will have better performance. However, you will not be able to use Resource Access Control Facility Version 1, Release 7 (RACF 1.7) to control authorization of ISMF functions. For more information on RACF 1.7, see the *RACF Security Administrator's Guide*.

Recovery: It is recommended that you put the new execution data sets in a place where your existing backup and recovery procedures will apply.

Making ISMF Available to TSO Users

Before TSO users can invoke ISMF, they must be allocated to the execution data sets. To make these execution libraries available to them, you must either modify an existing TSO LOGON procedure, create a new TSO LOGON procedure, or set up a CLIST that allocates these data sets. For information on how to accomplish each of these tasks, see *System Generation*.

ISPF Table Libraries

In addition to the execution data sets, you will wish to ensure that each ISMF user has a personal partitioned data set (PDS) available as a table output library. The ISMF SAVE command stores the saved list in this PDS using ISPF table services. This data set must be allocated with the DDNAME of ISPTABL; only one PDS may be used in this allocation because concatenated PDSs are not allowed on output. The data set name you choose might be of this form:

DSN=&userid.ISMF.DGTTABL

Be sure also to place this data set name in the concatenation for the ISPF table input library (ISPTLIB) so that the output of the SAVE command is available for subsequent use.

The DD statements for the table output library would then look like:

```
ISPTABL DD DSN=&userid.ISMF.DGTTABL,DISP=OLD
```

For the table input libraries, they would look like:

```
ISPTLIB DD DSN=&userid.ISMF.DGTTABL,DISP=SHR
         DD DSN=SYS1.DGTTLIB,DISP=SHR
         . . .
```

The table output library, *&userid.ISMF.DGTTABL*, has a fixed block record format and a record length of 80. The block size must be a multiple of 80 and be compatible with the other data sets in the concatenation.

Placement Considerations for the Load Modules

When you have completed sysgen processing, you will have several choices of where to place the load module libraries. The choice you make will affect the performance of your system. For example, there are maintenance considerations if the load modules exist anywhere other than SYS1.DGTTLIB. In this case, you must notify SMP of the location of these modules. For more information on performance issues, see the preceding section, "Where to Put the Execution Libraries" on page 24.

You can choose from one of the following:

1. Leave SYS1.DGTTLIB where it is and add the data set name to the LOGON JCL, or to a new or existing CLIST for use with TSO.
2. Leave SYS1.DGTTLIB where it is and add the data set name to the link library list (LNKLSTxx) or to the LPA library list (LPALSTxx) of the SYS1.PARMLIB. For information on setting up the link library list or the LPA library list, see *Initialization and Tuning*.
3. Make the load modules a part of either the system link library or the link pack area.

If you are not using SYS1.DGTTLIB at execution time and you have backed up and copied all the load modules, you can delete the SYS1.DGTTLIB data set.

For more information on placement, see *System Generation*, *DFDSS/ISMF Installation Planning Guide*, and *DFHSM with ISMF*.

Authorization Considerations

You can use RACF 1.7 to protect and report on ISMF. Using the program control feature, you can set up authorization levels for ISMF function: either all of ISMF, or individual line operators and commands.

Because program control allows you to determine who can execute which ISMF function, you can use it to set up an authorization scheme that applies to both individual users and user groups. If you plan to use RACF 1.7 to restrict ISMF functions, you must make the load module library a part of the system link library.

In addition to authorizing ISMF functions, you can use standard RACF authorization checking to limit access to individual data sets, volumes, or catalogs. Used in conjunction with program control, authorization checking will ensure that appropriate ISMF data and function are accessible to the users when needed. For detailed information on how to restrict ISMF functions, see *ISMF User's Guide*.

Chapter 4. Updating the Publications Library

When planning for installation of Data Facility Product, you should also plan to reorganize your system publications library for MVS/XA. Publications relevant to specific products should be ordered when those products are ordered, and added to the System Library Subscription Service (SLSS).

New manuals have been issued to support Data Facility Product to replace the manuals for Data Facility Device Support (DFDS) and Data Facility Extended Function (DFEF), as well as many of the OS/VS2 MVS publications. Other OS/VS2 MVS manuals apply also to the MVS/XA environment and should be retained for your reference. The new manuals, and the manuals that should be replaced, are discussed in the following section.

MVS/XA Data Facility Product Publications

The following table lists the publications that are specific to the MVS/XA Data Facility Product environment. The MVS/XA books in the left column replace the OS/VS, DFDS, DFEF, or previous DFP books in the right column.

Note: * indicates that this publication is classified as restricted material of IBM.

MVS/XA DFP Book	OS/VS, DFDS, DFEF, or Previous DFP Book(s)
* <i>MVS/XA Access Method Services Logic, Volumes 1 and 2, LY26-3953 and LY26-3997</i>	<i>OS/VS2 Access Method Services Logic, SY35-0010</i> <i>Data Facility Extended Function: Access Method Services Logic, LY26-3888</i>
* <i>MVS/XA BDAM Logic, LY26-3893</i>	<i>OS/VS2 BDAM Logic, SY26-3831</i>

MVS/XA DFP Book	OS/VS, DFDS, DFEF, or Previous DFP Book(s)
<i>MVS/XA Catalog Administration Guide, GC26-4138</i>	<i>Data Facility Extended Function: Access Method Services Reference, SC26-3967</i> <i>Data Facility Extended Function: Access Method Services Administration and Services, SC26-3966</i> <i>Planning for Enhanced VSAM Under OS/VS, GC26-3842</i> <i>OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838</i> <i>OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications, GC26-3819</i> <i>OS/VS2 MVS CVOL Processor, GC26-3864</i>
<i>*MVS/XA Catalog Diagnosis Reference, LY26-3956</i>	<i>Data Facility Extended Function: Catalog Diagnosis Reference, SY26-3886</i>
<i>*MVS/XA Checkpoint/Restart Supervisor Call Logic, LY26-3957</i>	<i>OS/VS2 Checkpoint/Restart Logic, SY26-3820</i>
<i>MVS/XA Checkpoint/Restart User's Guide, GC26-4139</i>	<i>OS/VS2 MVS Checkpoint/Restart, GC26-3877</i>
<i>*MVS/XA Common VTOC Access Facility Diagnosis Reference, LY26-3958</i>	<i>Data Facility Device Support: Common VTOC Access Facility Diagnosis Reference, SY26-3882</i>
<i>*MVS/XA CVOL Processor Logic, LY26-3895</i>	<i>OS/VS2 CVOL Processor Logic, SY26-3860</i>
<i>*MVS/XA DADSM Diagnosis Reference, LY26-3961</i>	<i>Data Facility Device Support: DADSM Diagnosis Reference, SY26-3881</i>
<i>MVS/XA Data Administration Guide, GC26-4140</i>	<i>OS/VS2 MVS Data Management Services Guide, GC26-3875</i>
<i>MVS/XA Data Administration: Macro Instruction Reference, GC26-4141</i>	<i>OS/VS2 MVS Data Management Macro Instructions, GC26-3873</i>
<i>MVS/XA Data Administration: Utilities, GC26-4150</i>	<i>OS/VS2 MVS Utilities, GC26-3902</i>
<i>MVS/XA Data Facility Product Version 2: Customization, GC26-4267</i>	<i>MVS/XA Interactive Storage Management Facility Customization Guide, GC26-4267</i>

MVS/XA DFP Book	OS/VS, DFDS, DFEF, or Previous DFP Book(s)
<p><i>*MVS/XA Data Facility Product Version 2: Diagnosis Guide, LY27-9521</i></p>	<p><i>*MVS/XA Catalog Diagnosis Guide, LY26-3955</i></p> <p><i>Data Facility Extended Function: Catalog Diagnosis Guide SY26-3887</i></p> <p><i>*MVS/XA DADSM and Common VTOC Access Facility Diagnosis Guide, LY26-3960</i></p> <p><i>Data Facility Device Support: DADSM and Common VTOC Access Facility Diagnosis Guide, SY26-3880</i></p> <p><i>*MVS/XA Data Facility Product Common Services Diagnosis Guide, LY26-3959</i></p> <p><i>*MVS/XA Interactive Storage Management Facility Diagnosis Guide, LY26-3906</i></p> <p><i>*MVS/XA Media Manager Diagnosis Guide and Reference, LY26-3965</i></p> <p><i>Data Facility Device Support: OS/VS2 Media Manager Diagnosis Guide and Reference, SY26-3884</i></p>
<p><i>MVS/XA Data Facility Product Version 2: Diagnosis Reference, LY27-9530</i></p>	<p>"Diagnostic Aids" sections from many logic manuals and diagnostic reference manuals contributed to this book</p>
<p><i>MVS/XA Data Facility Product Version 2: General Information, GC26-4142</i></p>	<p>None</p>
<p><i>MVS/XA Data Facility Product Version 2: Master Index, GC26-4146</i></p>	<p>None</p>
<p><i>MVS/XA Data Facility Product Version 2: Planning Guide, GC26-4147</i></p>	<p><i>Data Facility Extended Function: Access Method Services Administration and Services, SC26-3966</i></p> <p><i>Data Facility Device Support: User's Guide and Reference, SC26-3952</i></p>
<p><i>MVS/XA Installation: System Generation, GC26-4148</i></p>	<p><i>OS/VS2 System Programming Library: System Generation Reference, GC26-3792</i></p>

MVS/XA DFP Book	OS/VS, DFDS, DFEF, or Previous DFP Book(s)
<i>MVS/XA Integrated Catalog Administration: Access Method Services Reference, GC26-4135</i>	<i>OS/VS2 Access Method Services, GC26-3841</i> <i>OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option, SC26-3916</i> <i>Data Facility Extended Function: Access Method Services Reference, SC26-3967</i>
<i>MVS/XA Integrated Catalog Administration: Access Method Services Reference Summary, GX26-3724</i>	None
<i>MVS/XA Interactive Storage Management Facility User's Guide, GC26-4266</i>	None
<i>*MVS/XA ISAM Logic, LY26-3894</i>	<i>OS/VS2 ISAM Logic, SY26-3833</i>
<i>MVS/XA Linkage Editor and Loader User's Guide, GC26-4143</i>	<i>OS/VS Linkage Editor and Loader, GC26-3813</i>
<i>*MVS/XA Linkage Editor Logic, LY26-3963</i>	<i>OS/VS Linkage Editor Logic, SY26-3815</i>
<i>*MVS/XA Loader Logic, LY26-3901</i>	<i>OS/VS Loader Logic, SY26-3814</i>
<i>MVS/XA Magnetic Tape Labels and File Structure Administration, GC26-4145</i>	<i>OS/VS Tape Labels, GC26-3795</i>
<i>*MVS/XA Media Manager Diagnosis Reference, LY26-3965¹</i>	<i>*MVS/XA Media Manager Diagnosis Guide and Reference, LY26-3965</i> <i>Data Facility Device Support: OS/VS2 Media Manager Diagnosis Guide and Reference, SY26-3884</i>
<i>*MVS/XA Open/Close/EOV Logic, LY26-3966</i>	<i>OS/VS2 Open/Close/EOV Logic, SY26-3827</i>
<i>*MVS/XA SAM Logic, LY26-3967</i>	<i>OS/VS2 SAM Logic, SY26-3832</i> <i>*OS/VS2 MVS SAM-E Logic, LY26-3855</i>
<i>MVS/XA System – Data Administration, GC26-4149</i>	<i>OS/VS2 System Programming Library: Data Management, GC26-3830</i> <i>Data Facility Device Support: User's Guide and Reference, SC26-3952</i>

¹ The diagnosis guide information from this publication has been transferred to the new Diagnosis publication. Consequently, this publication has been retitled and is considered redundant to the DFP diagnosis task.

MVS/XA DFP Book	OS/VS, DFDS, DFEF, or Previous DFP Book(s)
<i>*MVS/XA Utilities Logic, LY26-3968</i>	<i>OS/VS Utilities Logic, SY35-0005</i>
<i>*MVS/XA VIO Logic, LY26-3900</i>	<i>OS/VS2 VIO Logic, SY26-3834</i>
<i>MVS/XA VSAM Administration Guide, GC26-4151</i>	<i>OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838</i> <i>OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications, GC26-3819</i> <i>Planning for Enhanced VSAM Under OS/VS, GC26-3842</i> <i>Data Facility Extended Function: Access Method Services Administration and Services, SC26-3966</i> <i>OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option, SC26-3916</i>
<i>MVS/XA VSAM Administration: Macro Instruction Reference, GC26-4152</i>	<i>OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838</i> <i>OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications, GC26-3819</i>
<i>MVS/XA VSAM Catalog Administration: Access Method Services Reference, GC26-4136</i>	<i>OS/VS2 Access Method Services, GC26-3841</i> <i>OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option, SC26-3916</i> <i>Data Facility Extended Function: Access Method Services Reference, SC26-3967</i>
<i>*MVS/XA VSAM Logic, LY26-3970</i>	<i>OS/VS2 VSAM Logic, SY26-3825</i>

Changes in the MVS/XA Data Facility Product Library

Our objective for the MVS/XA Data Facility Product library is to present the information so that users can easily find what they need to know in order to accomplish particular tasks. These tasks change somewhat as the product changes.

For detailed information about the changes and additions in each book that was updated for this release of Data Facility Product, see the table of updates in the introduction to the *Data Facility Product: Master Index*.

Customization Restructure

With the introduction of the interactive storage management facility (ISMF), we introduced the *MVS/XA Interactive Storage Management Facility Customization Guide*. Now we are enlarging this book to include customization information for all of Data Facility Product. The new *Data Facility Product: Customization* replaces the *ISMF Customization Guide* and also includes information from the following books:

- *Data Administration Guide*: Chapter 7, "DCB Exit Routines"
- *Data Administration: Macro Instruction Reference*: Appendix A, "Status Information Following an Input/Output Operation"
- *Data Administration: Utilities*: Appendix A, "Exit Routine Linkage," and Appendix C, "Processing User Labels"
- *Magnetic Tape Labels and File Structure Administration*: Chapter 4, "Nonstandard Labels"; Chapter 6, "Volume Label Verification and Volume Label Editor Routines"; and Appendix D, "Version 3 Installation Exits"
- *System—Data Administration*: Chapter 2 section entitled "Appendages"; Chapter 5, "Exit Routines"; Chapter 9 "CATALOG, SCRATCH, and RENAME Dummy Modules"; and Appendix D, "Example of an OPEN Installation Exit Module"
- *VSAM Administration Guide*: Chapter 11, "User-Written Exit Routines," and Appendix G, "Datestamp Routine"

Data Facility Product: Customization is intended to be used primarily by system programmers, but application programmers may need the information it contains on DCB exit routines and on VSAM exits.

Diagnosis Restructure

There are two new Diagnosis publications, *MVS/XA Data Facility Product Version 2: Diagnosis Guide* and *MVS/XA Data Facility Product Version 2: Diagnosis Reference*.

The *Diagnosis Guide* contains the information you will need to diagnose problems in any DFP product component. It supersedes the following diagnosis guides:

- *Catalog Diagnosis Guide*, LY26-3955
- *Common Services Diagnosis Guide*, LY26-3959
- *DADSM and Common VTOC Access Facility Diagnosis Guide*, LY26-3960
- *Interactive Storage Management Facility Diagnosis Guide*, LY26-3906

- The diagnosis guide portion of *Media Manager Diagnosis Guide and Reference*, LY26-3965

The *Diagnosis Reference* contains the diagnosis tools and service aid information required to support the diagnosis tasks described in the *Diagnosis Guide* publication. This information is taken from existing DFP diagnosis reference and program logic publications. Although these publications remain in the DFP library, they are now redundant to the diagnosis task.

MVS/XA Storage Management Library

The MVS/XA Storage Management Library (SML) provides guidance in establishing efficient storage management practices and in moving toward system-managed storage. The library is designed to work in concert with the libraries that document the various storage-related hardware and software products. The SML provides a framework for determining when information in another library is relevant to a particular task.

- *Configuring Storage Subsystems*, GC26-4262
- *Focus on Storage Management*, GC26-4260
- *Leading an Effective Storage Administration Group*, GC26-4261
- *Managing Data Sets*, GC26-4263
- *Managing Storage Pools*, GC26-4264
- *Storage Management Reader's Guide*, GC26-4265

Publications for the Storage Management Family of Products

The storage management family of products, which includes Data Facility Product, performs tasks essential to storage management. These products are designed to work together so that the total benefits derived from the family may exceed the sum of benefits of each individual product.

You may want to refer to the *Storage Management Reader's Guide* for a complete listing of the books in each library.

Data Facility Data Set Services (DFDSS)

The Data Facility Data Set Services (DFDSS) library provides guidance in copying or moving data sets between volumes of like or unlike device type, dumping and restoring data sets (as well as full or partial volumes), and reducing or eliminating DASD free-space fragmentation.

- *Data Facility Data Set Services: General Information*, GC26-4123
- *Data Facility Data Set Services: User's Guide and Reference*, SC26-4125
- *Data Facility Data Set Services/Interactive Storage Management Facility Installation Planning Guide*, SC26-4129

Data Facility Hierarchical Storage Manager (DFHSM)

The Data Facility Hierarchical Storage Manager (DFHSM) library provides guidance in managing data availability and space maintenance.

- *Data Facility Hierarchical Storage Manager General Information*, GH35-0092
- *Data Facility Hierarchical Storage Manager: System Programmer's Guide*, SH35-0085
- *Data Facility Hierarchical Storage Manager Version 2.2.1 with the Interactive Storage Management Facility*, SH35-0108

Resource Access Control Facility (RACF)

The Resource Access Control Facility (RACF) library provides information about establishing resource and system data security.

- *Resource Access Control Facility (RACF): General Information Manual*, GC28-0722
- *Resource Access Control Facility Security Administrator's Guide*, SC28-1340

Data Facility Sort (DFSORT)

The Data Facility Sort (DFSORT) library provides guidance for sorting, merging, and copying data in MVS and MVS/XA environments. In addition, it explains how to alter, delete, or insert records as they are being sorted, merged, or copied.

- *DFSORT General Information*, GC33-4033

Storage Subsystem Library

The Storage Subsystem Library (SSL) provides hardware and software information about storage management for direct access storage devices (DASD), magnetic tape units, and storage control units (with and without cache).

In addition to describing characteristics, features, capabilities, and configuration options for specific storage devices, the SSL specifies software requirements and options. The library also provides detailed instructions for effective planning, installation, and use of the devices in various operating environments.

Device-related programming enhancements are documented first in the SSL. These enhancements are later incorporated into the next release of the appropriate software library.

For a list of the manuals in the SSL, see the *Storage Management Reader's Guide*.

Other Product-Related Publications

The manuals listed below may be of interest if you plan to install or use the appropriate licensed program.

Prerequisite or Corequisite Licensed Programs

MVS/System Product JES2 or JES3

- *MVS/System Product Version 2 General Information Manual*, GC28-1118
- *MVS/Extended Architecture JES3 Introduction*, GC23-0049

Assembler H Version 2

- *Assembler H Version 2: General Information*, GC26-4035

System Modification Program (SMP or SMP/E)

- *System Modification Program (SMP) System Programmer's Guide*, GC28-0673
- *System Modification Program Extended General Information Manual*, GC28-1106
- *System Modification Program Extended User's Guide*, SC28-1302

Environmental Recording, Editing, and Printing (EREP)

- *Environmental Recording, Editing, and Printing (EREP) User's Guide and Reference*, GC28-1178

Device Support Facilities

- *Device Support Facilities User's Guide and Reference*, GC35-0033

Interactive System Productivity Facility (ISPF or ISPF/PDF)

- *Interactive System Productivity Facility (ISPF) and ISPF/Program Development Facility Version 2 Release 2 General Information Manual*, GC34-4041

Time Sharing Option (TSO or TSO/E)

- *MVS/Extended Architecture System Programming Library: Time Sharing Option*, GC28-1173

Optional Licensed Programs

- *Advanced Communications Function for TCAM Version 2 General Information: Introduction, GC30-3057*
- *Advanced Communications Function for VTAM Version 2 General Information, GC27-0608*
- *CICS/OS/VS General Information, GC33-0155*
- *IMS/VS Version 2 General Information Manual, GC26-4180*
- *OS/VS BTAM, GC27-6980 (with BTAM/SP supplement, SC27-0604)*
- *OS/VS1 and OS/VS2 MVS Cryptographic Unit Support General Information, GC28-1015*
- *OS/VS1 and OS/VS2 MVS Cryptographic Unit Support Installation Reference, SC28-1016*
- *OS/VS1 and OS/VS2 MVS Programmed Cryptographic Facility General Information, GC28-0942*
- *OS/VS1 and OS/VS2 MVS Programmed Cryptographic Facility Installation Reference, SC28-0956*
- *Resource Measurement Facility (RMF) Version 3 General Information, GC28-1115*

Part 2. Planning for Integrated Catalog Facility Catalogs



Chapter 5. Integrated Catalog Facility Catalog Overview

Integrated catalog facility catalogs functionally replace VSAM catalogs and OS CVOLs. An integrated catalog facility catalog can be the master catalog or a user catalog. Integrated catalog facility catalogs, VSAM catalogs, and OS CVOLs can coexist and can be cataloged in an integrated catalog facility catalog.

An installation can install integrated catalog facility catalogs or migrate to them as fast or as slowly as it chooses. All existing catalogs (VSAM and OS CVOLs) continue to operate in the same way. In general, the following need not be changed when installing or migrating to integrated catalog facility catalogs:

- Access method services job streams
- User job streams that continue to use existing catalogs
- Programs (except those dependent on the format of data set names in the VTOC)

Programs that rely on tracing catalog control block chains and pointers in the common service area (CSA) and pageable link pack area (PLPA) must be updated to address the new catalog address space (CAS). For additional information about integrated catalog facility catalogs and CAS, see *Catalog Administration Guide*.

The Integrated Catalog Facility Catalog

An integrated catalog facility catalog consists of two types of components: the basic catalog structure (BCS) and one or more VSAM volume data sets (VVDS). Catalog information about VSAM data sets is split between the BCS and the VVDS. Information about non-VSAM data sets is contained only in the BCS.

Changes caused by loading, updating, or extending VSAM data sets result in changes to the VVDS; the BCS information remains unchanged.

Non-VSAM data sets are described entirely in the BCS and not in the VVDS.

Figure 4 on page 40 shows the relationship of the BCS and the VVDS.




Figure 4. Relationship of the BCS and the VVDS

The BCS

The BCS is a key-sequenced data set (KSDS). It contains volume, data set security, ownership, and associated information for VSAM and non-VSAM data sets.

A BCS may point to VVDSs on any number of volumes. A BCS may also point to a volume on which only non-VSAM and generation data group data sets reside. The non-VSAM and generation data group data set information required by the integrated catalog facility catalog is contained in the BCS itself, and the VVDS does not contain the non-VSAM data set information.

Related information in the BCS is grouped into variable-length, spanned records. The BCS uses keys that are the data set names (plus one character for extensions). A control interval may contain multiple BCS records. Logically related data is consolidated in the BCS to reduce the number of I/Os necessary for catalog processing. For further information about key-sequenced data sets, see *VSAM Administration Guide*.

The VVDS

The VVDS is an entry-sequenced data set (ESDS) that has 4K-byte (K equals 1024) control intervals. It contains the information about VSAM data sets residing on the volume with the VVDS. A VVDS may have VSAM data sets cataloged in up to 36 integrated catalog facility catalogs. The VVDS is composed of a minimum of two records: one VSAM volume control record (VVCR) and one (or more) VSAM volume record(s) (VVR).

The first logical record in a VVDS is the VVCR. The VVCR contains information on managing the VVDS DASD space; it also contains entries for up to 36 integrated catalog facility catalogs that have BCS or VSAM data set components on that volume.

The remaining logical records in the VVDS are VVRs and contain data set characteristics: for example, extent information or frequently used relative byte addresses (RBAs) for a BCS or for the VSAM data sets residing on that volume.

A VVDS may be implicitly or explicitly defined. The VVDS is usually dynamically created (implicitly defined), using default primary and secondary space allocation quantities the first time a catalog or a VSAM data set is allocated on a volume. Specific space allocation quantities may be explicitly specified when defining a VVDS. For further information about defining the VVDS, see *Catalog Administration Guide*.

A VVDS is recognized by the VVDS data set name `SYS1.VVDS.Vvolser`, in which *volser* is the volume serial number of the volume in which the VVDS resides.

All DASD space management is provided by DADSM, and each data set cataloged in an integrated catalog facility catalog has corresponding VTOC DSCBs. Two new flags exist in the Format-1 DSCB in the OPTCD field. OPTCD of X'80' indicates the data set is defined in an integrated catalog facility catalog (for example, a VVDS), and OPTCD of X'C0' indicates an integrated catalog facility catalog BCS entry (either the data or index component).

Because of its special use, a VVDS cannot be password protected, it cannot be exported or imported, and its attributes cannot be altered.

Entries in a Catalog

A catalog contains entries that describe VSAM and/or non-VSAM objects. One or more catalog entries are built to describe the defined objects. The objects you may define are:

- *User catalog*, a collection of information about non-VSAM and VSAM objects
- *Cluster*, a collection of data organized in key-sequence, entry-sequence, linear, or relative record
- *Alternate index*, an index that is organized in key sequence and is related to a key-sequence or entry-sequence cluster
- *Path*, a name for the combination of an alternate index and its base cluster or an alias for a VSAM data set
- *Page space*, an amount of DASD space formatted for the operating system for paging operations

The non-VSAM objects you may define are:

- *Non-VSAM data set*, a sequential, partitioned, direct, or indexed-sequential data organization (not VSAM data organization)

- *Generation data group*, a collection of non-VSAM data sets grouped together in a time-dependent manner
- *Alias*, an alternative name for a user catalog or non-VSAM data set

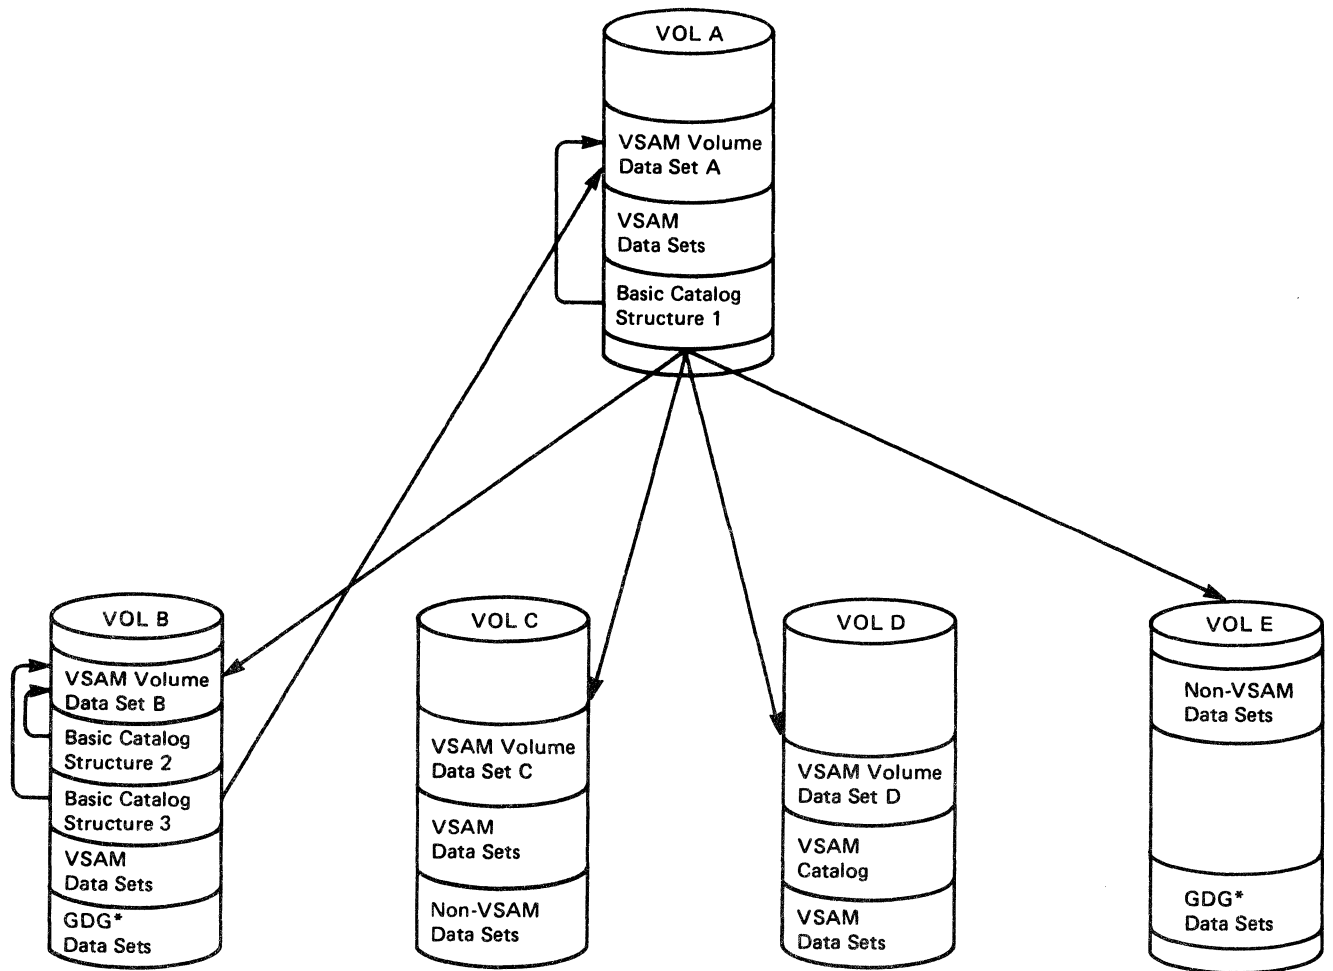
Figure 5 on page 43 shows the relationship of an integrated catalog facility catalog and its objects.

For further information on defining objects, see *Catalog Administration Guide* or *VSAM Administration Guide*.

Alternate Master Catalog

The alternate master catalog provides an easy way to back up the master catalog in case of damage to the master catalog. It also provides a simple method of converting a VSAM master catalog to an integrated catalog facility master catalog without the need for an additional system.

Several steps are necessary to generate an alternate master catalog. An alternate master catalog may be specified at IPL (initial program load) time. The operator specifies which SYSCATnn member of the SYS1.NUCLEUS data set is to be used to find the master catalog during the current IPL. The selected master catalog must already exist and have the necessary system data sets cataloged in it. For further information, see *Catalog Administration Guide*.



	has cataloged entries on:
Basic Catalog Structure 1	VOL A, VOL B, VOL C, VOL D, VOL E
Basic Catalog Structure 2	VOL B
Basic Catalog Structure 3	VOL A, VOL B

*Generation Data Group Data Sets

Figure 5. Integrated Catalog Facility Catalog Relationships



Chapter 6. Integrated Catalog Facility Catalog Installation

Most installations depend on the availability of cataloging facilities to run production job streams and to support online users. Modifying these facilities to use integrated catalog facility catalogs requires thorough planning, beginning with an analysis of current use of catalogs.

Changing the catalog environment may involve a revision of conventions and catalog administration responsibilities. New procedures for catalog backup and recovery and for reorganization may be needed. Some users' job streams might be affected. For example, if EXPORTRA is used for data set backup, it will result in an error condition; EXPORTRA is not supported for an integrated catalog facility catalog.

The extent of the revision depends on the current status of the catalog environment. In a well-designed and carefully controlled catalog environment, the replacement of existing catalogs with integrated catalog facility catalogs should be straightforward and have little impact on the installation. However, if your current catalog environment has limitations because of prior use of catalog facilities (for example, OS CVOLs, or unusual user catalogs), migration may take longer.

An Installation Planning Process

The steps listed below are suggested procedures to follow for installation planning:

1. List the objectives for the catalog environment and state specific requirements in the order of their priority.
2. Examine the objectives and conflicting and complementary relationships.
3. Compare the current catalog environment with the objectives to determine the necessary tasks to support the objectives.
4. List the specific installation tasks.
5. Organize the tasks into an installation plan, establishing an order of precedence and preparing a schedule for the conversion by location, by system, or by application grouping.
6. Assign responsibilities to personnel.
7. Execute the installation plan.
8. Assign responsibilities for monitoring the catalog environment.

The Catalog Environment and the Installation Tasks

Listed below are some elements of the catalog environment to consider when installing and maintaining cataloging facilities. The list, although not exhaustive, is intended to provide a basis for discussion and for creating a list of tasks or procedures for a specific installation.

- Conventions for data sets and catalogs
- The configuration of catalogs
- Individual catalog structures
- Procedures for data sets and catalogs

The following sections present guidelines within these areas for anyone installing integrated catalog facility catalogs or migrating from VSAM catalogs and OS CVOLs to integrated catalog facility catalogs. If you are planning changes in addition to those required for integrated catalog facility catalogs, there may be other tasks that are not discussed in this manual.

Conventions for Data Sets and Catalogs

For maximum reliability and efficiency, all catalogs should be integrated catalog facility catalogs and all permanent data sets should be cataloged.

Data set naming conventions generally should require all names to have multiple levels. The Resource Access Control Facility (RACF) program product (5740-XXH) protects only data sets with multilevel names. Data sets with names "A" and "A.B" cannot both be cataloged in the same catalog. (However, it is possible to have "A.B" and "A.B.C" in the same catalog.) The highest-level qualifier should identify the owning user or application.

If your installation has locally written programs that are dependent on the format of data set names in the VTOC, the programs should be revised to allow for the changed data set names that occur in the integrated catalog facility catalog environment. For information about the format for these names, see "VTOC Entries for Integrated Catalog Facility Catalogs" in *Catalog Administration Guide*.

JOBCAT and STEPCAT statements should be used only with jobs that require them: for example, jobs that process catalogs as data sets. For further information about JOBCAT and STEPCAT statements, see *Access Method Services Reference*.

A selected alternate master catalog should be created explicitly, and should have the necessary system data sets cataloged in it.

A catalog coordinator should control the catalog configuration, designate catalog owners, and establish guidelines for catalog construction and procedures.

Each catalog should have a designated owner who is responsible for achieving the objectives set by the installation for security, performance, availability, and for user assistance.

The owner should be responsible for the performance of the catalog. Performance options are controlled by parameters of the DEFINE and ALTER commands, so, generally, the owner creates and maintains the catalog. For further information about these parameters, see *Access Method Services Reference*.

One of the most significant factors influencing catalog performance is whether the catalog is to be shared between systems.

If multiple users or applications are to share a catalog, the possible impact of a catalog failure should be considered. It may be desirable to establish separate catalogs for critical applications. Each catalog must have tested and documented backup and recovery procedures. Ensuring that such procedures exist is another responsibility of the catalog owner. See *Catalog Administration Guide* for information about these procedures.

To make orderly changes in catalog access and configuration, the catalog owner should maintain a current access list and configuration diagram for each catalog.

Conventions Summary

With the preceding factors in mind, consider the following general guidelines for conventions or standards.

- A single catalog coordinator should control the installation's configuration of catalogs and establish guidelines for catalog structures and procedures.
- Each catalog should have an owner who is solely responsible for the management and maintenance of that catalog.
- All catalogs should be password protected.
- All catalogs should have tested and documented backup and recovery procedures.
- All permanent data sets should be cataloged.
- All data sets should have multilevel names, with the highest level identifying the user or application.
- JOBCAT and STEPCAT statements should be discouraged (except as required for certain catalog management functions).
- The master catalog should have an alternate master catalog to use as a backup.

Catalog Configuration

The configuration of catalogs in an installation is the set of all user and master catalogs for all connected (via shared DASD) systems.

Figure 6 on page 49 is a configuration diagram. It shows the catalog names and volumes, and connections from the master catalog and alternate master catalog via aliases. The diagram contains several catalogs, including:

- A master catalog

The master catalog, MAS.CAT, is on volume SY1VOL and connects to all the user catalogs, including the alternate master catalog with aliases.

- An alternate master catalog

The alternate master catalog, ALT.CAT, is on volume SY2VOL and connects (dotted lines) to the master catalog as well as to the user catalogs. The alternate master catalog provides a backup if the master catalog should be damaged.

- User catalogs

All the other catalogs in the configuration are for users. For example, CAT.ABC is on VOL.123 and is owned by T. Smith. The catalog is connected to the master catalog by the aliases INVEN, MANU, and SALES.

Each catalog on your diagram should contain the catalog name, the volume the catalog is on, the owner's name, and the aliases. The diagrams can be used to depict the current as well as the intended configuration, and are an important tool for planning the installation of integrated catalog facility catalogs. As each new catalog is added, or when existing catalogs are reconstructed, the configuration diagram should be updated.

The user should not be concerned about which volumes are catalog volumes. This is usually a factor only when MSS volumes are considered for catalog residence.

If the MSS environment is JES3, the catalog virtual volume must be mounted before it is referenced. Instead, the installation should mount the necessary virtual volumes either by using mount commands or by including catalog volumes in the volume attribute list (VATLIST). A JOBCAT and STEPCAT can be used to mount the volume; however, as described above, this is not normally a desirable solution.

If the MSS environment is JES2, reference to a catalog causes the virtual volume to be mounted. Catalogs always have the BIND attribute; so, on a virtual volume, a catalog occupies staging space as long as it is open.

In general, catalogs should not reside on virtual volumes in a JES2 or JES3 environment.

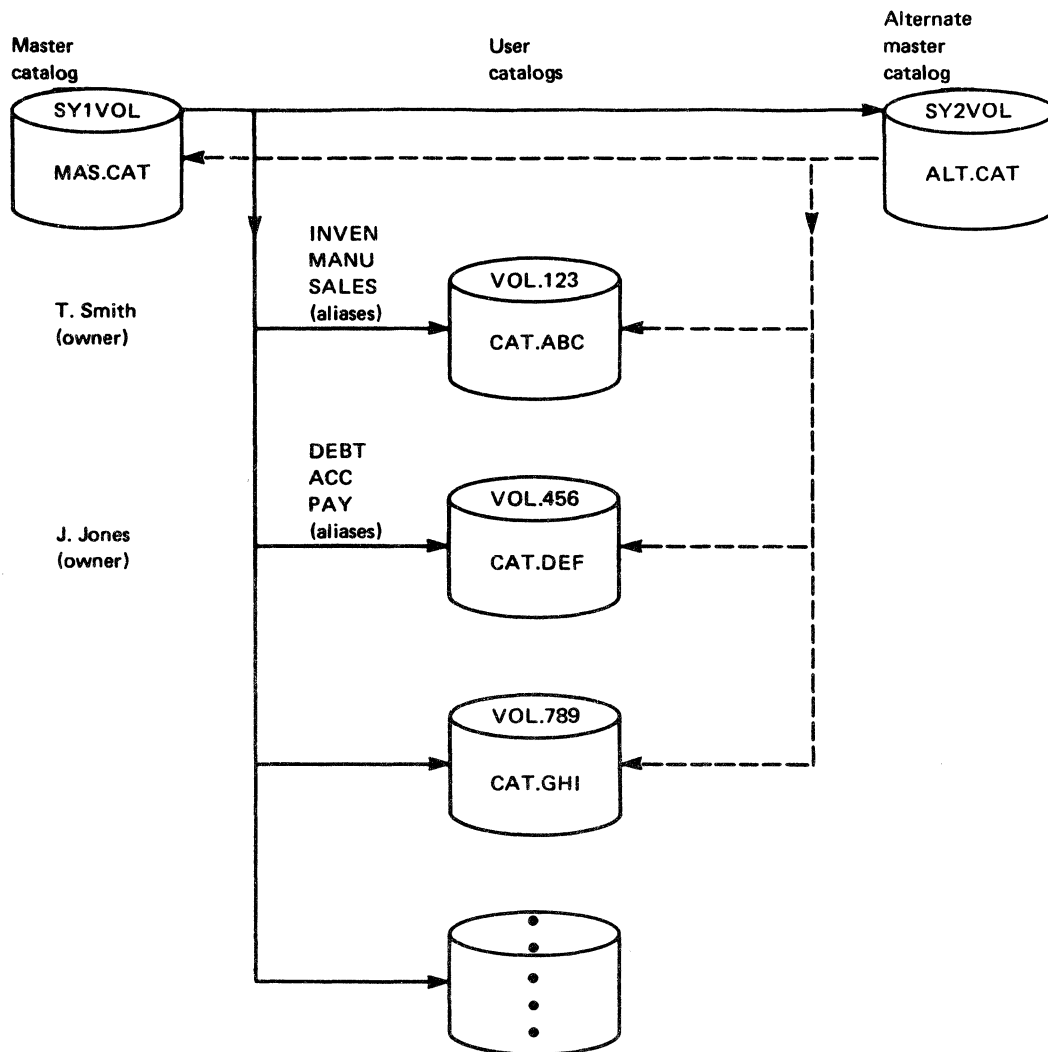


Figure 6. Catalog Configuration Diagram

The configuration should include an alternate master catalog to enable you to recover a damaged master catalog. If the installation has multiple systems, the damaged catalog might be recovered as a user catalog from another system. Otherwise, a true alternate master catalog with its own page spaces and entries should be available.

The number of catalogs is determined by the availability requirements, the need to isolate critical applications, and the time required for recovery.

The time required for recovery is also a consideration. If a damaged catalog contained non-VSAM entries for many volumes, processing the VTOCs or system management facilities (SMF) data may require longer than the tolerable recovery time. This is true for VSAM data sets in terms of the time required to run DIAGNOSE (to validate data structures to detect structure errors) and to perform the indicated recovery actions. Configuration planning should include a "failure analysis" for critical catalogs to determine which applications and/or users are affected and how long a recovery might take.

If multiple integrated catalog facility catalogs are used for availability, then two integrated catalog facility catalogs should not reside on the same volume. If they do, they share the VVDS on that volume, and are susceptible to the loss of the same VVDS and the loss of the volume.

Catalogs with high availability needs should be placed on units with alternate path access (at least from the most critical processor).

It is necessary to consider both the current and the target configurations when planning the installation or conversion. It may also be desirable to pass through some interim configurations. Generally, it is easier to convert catalogs on a one-for-one basis and immediately reconfigure with the split and merge facilities.

Catalog Configuration Summary

Based on the above considerations, the following general guidelines may be used for designing the catalog configuration:

- The master catalog should contain only the required entries: system data sets, user catalog connectors, and aliases. It should be protected from update.
- For data sets with high security requirements, provide a separate catalog or catalogs for maximum protection from a holder of RACF alter authority.
- To simplify recovery, do not allow two catalogs of any type to reside on the same volume.
- To reduce the impact of a catalog failure, provide each major application or user group with a separate catalog.
- Consider the number of data sets and data volumes to be matched against a restored catalog, and limit the scope of each catalog so the recovery time is within acceptable limits.
- Subject to performance and space utilization considerations, group data volumes by application or user group corresponding to a single catalog. When possible, catalog all data sets on a volume in the same catalog.
- When selecting the catalog unit, consider the contention caused by usage of the device and its path(s) from all systems.
- Try to position the catalog so that sharing by multiple systems is not required (or is minimized).
- If the rate of catalog requests to a given catalog is expected to be very high, attention should be given to performance considerations. It may be desirable to split such a catalog for better performance.
- Maintain a current catalog configuration diagram.

Catalog Construction

The construction of each catalog is determined by the parameters specified at the time the catalog is defined. Certain parameters not affecting the way the catalog data is physically stored may be modified by the ALTER command. Generally, it is necessary to understand the objectives and the intended use of a catalog before it can be properly constructed.

Every catalog for the new environment should be defined as an integrated catalog facility catalog; the parameter used to define the integrated catalog facility catalog must be explicitly coded because the default is a VSAM catalog. Alias pointers for the catalog must be established in each system that will try to access it.

As each catalog is created, the impact of a catalog failure should be considered. Some applications may have so many availability requirements that it is undesirable to have them share a catalog with other, less critical applications. If access from more than one host is required, the sharing options must be specified. Finally, in selecting a volume for a catalog, the relationship between the catalog volume and the data sets is important to the design of the recovery procedure. Only those volumes that contain cataloged data sets need to be considered (by DIAGNOSE or VTOC inspection) during catalog recovery.

The owner (usually the creator) of each catalog should be identified in the catalog, at the time it is defined, by specifying the OWNER keyword. The owner is responsible for the security, performance, and availability of the catalog.

Catalog Construction Summary

The following guidelines, which are more than a summary of the points above, include performance considerations.

- Always specify the owner and the retention period.
- Always protect the catalog. If RACF is used, specify passwords to provide protection on non-RACF systems.
- Control the size and placement of the VVDS, the BCS data component, and the BCS index component. The VVDS and the BCS index may share a cylinder that is adjacent to the BCS data component.
- Embed the BCS index sequence set. This is the default.
- Replicate the BCS index. This is not the default.
- Specify SHAREOPTIONS (3 3) only if the catalog is never to be shared. VSAM code does not control this option.
- Allow some free space for future insertions.
- Set the control interval size for both the data and the index components of the BCS to 1024 bytes. This represents a compromise between conserving buffer space and avoiding spanned records.

- Control the control area size by specifying the secondary allocation catalog in tracks to provide consistency with the control interval size specifications.
- Let BUFNI and BUFND default to STRNO + 1. These numbers should be adequate for performance and still not require excessive space in the common service area (CSA) for buffers.
- Start with the default of STRNO = 2, then watch for enqueue waits on SYSZRPLW.catalogname using RMF (Resource Measurement Facility), or a similar facility. Frequent waits on this name do not necessarily indicate that the string number is too small. The I/O service times for catalog requests should be inspected to determine whether there are delays caused by device busy or path busy conditions (see "Setting Up the BCS" on page 67).

Procedures for Data Sets

The procedures for referencing and cataloging non-VSAM data sets automatically (that is, through JCL or dynamic allocation) do not depend on the type of catalog in use.

Certain procedures for VSAM data sets are, however, affected by the integrated catalog facility catalog. These are discussed in the next sections as they relate to the objectives for the catalog environment.

The areas most significantly impacted for VSAM data sets are space allocation and management. Integrated catalog facility catalogs have no volume ownership and no suballocated spaces.

Integrated catalog facility catalogs allow the allocation of VSAM data sets on any specified volume. In addition, the user may identify data sets in the VTOC. Because the data and index component names appear in the VTOC, it may be desirable to name these components explicitly. Each component (including each key range) of a VSAM data set is uniquely represented by a Format-1 DSCB in the VTOC. Key range data sets may have more than one key range per volume. A key range qualifier is added to the data set name in the DSCB to prevent a duplicate name condition arising.

The user can determine from the VTOC (or VTOC index) if sufficient space exists on a volume for a new allocation or extension. Only VSAM data sets may have as many as 119 to 123 extents per volume (DADSM uses 1 to 5 extents). Non-VSAM data sets may have only as many as 16 extents per volume.

Because the allocation information resides in the VVDS on the volume with the data, the volume must be mounted any time the VVDS is to be accessed to retrieve such information. If a DD statement is not supplied and the volume meets dynamic allocation requirements, the volume will be dynamically allocated. For example, if the user requires allocation information about VSAM data sets on unmounted volumes, a DD statement should be supplied in the LISTCATALOG procedure to cause the data volume(s) to be mounted. If the allocation information is not really required, the scope of the LISTCATALOG should be restricted.

If JOBCAT and STEPCAT statements are currently in use and will not be used in the new environment, they should be removed.

VSAM data set procedures usually do not influence catalog performance. An exception is the processing of data set extensions. Because all allocation (including additional extents) is handled by DADSM, catalog management does not have to construct the SMF type 69 records for VSAM volumes.

Backup and recovery procedures for VSAM data sets that rely on CRAs (catalog recovery area) must be changed for integrated catalog facility catalogs. As with non-VSAM data sets, VSAM data sets can be backed up individually or by volume; however, unlike non-VSAM data sets, there is no facility for recovering individual VSAM data sets from a volume dump. If this technique is used, the entire volume must be restored and all data sets on the volume then made current.

A step in the recovery of a BCS may be the use of the DEFINE command and keyword RECATALOG for individual data sets. The original define statements for a VSAM data set may be saved for use in such a situation; otherwise, input for the DEFINE RECATALOG command must be created at recovery time.

VSAM data sets can be placed (or relocated) independent of their catalog. This gives you increased flexibility, but you must plan recovery procedures carefully.

Nonspecific volume allocation is not supported for VSAM data sets. You must still identify in the DEFINE command the volume on which the VSAM data set is to be allocated.

With integrated catalog facility catalogs, some of the information necessary for VSAM OPEN is contained in the VVDS on the data volume(s). All volumes of a multivolume VSAM data set must be mounted for OPEN to complete.

Data Set Procedures Summary

The following items summarize the data set procedures required by or made desirable by integrated catalog facility catalogs.

- No changes are identified for non-VSAM data sets.
- A VVDS may have VSAM data sets cataloged in up to 36 integrated catalog facility catalogs.
- Volume ownership cannot be used to control VSAM data set allocation. If required, such control can be exercised by installation-written DADSM exits.
- All VSAM components cataloged in an integrated catalog facility catalog are allocated as distinct entities by DADSM. The target volume(s) for DEFINE operations must contain available space and available DSCBs. Specify names for data and index components as they will appear in the VTOC.
- The unit of DASD allocation for VSAM data may be in cylinders, records, or tracks. However, if the allocation is specified in tracks, DASD allocation will be in contiguous tracks.

- Unique allocations are not rounded to a cylinder or cylinders. Specify the space actually required.
- VSAM data sets may have as many as 119 to 123 extents (DADSM uses 1 to 5 extents per volume), but the extent limit for non-VSAM data sets is 16.
- Reusable VSAM data sets may not be suballocated.
- VSAM data sets may have the REUSE parameter, which allows you to use VSAM data sets as work files.
- Secondary extents on the first volume are not released when reusable data sets are opened for REUSE. Secondary extents on any other volume than the first volume will be released.
- Secondary extents for VSAM data sets are not released when such data sets are opened for reuse.
- Allocation information for VSAM data sets is stored in the VVDS on the data volume(s). When listing catalog entries, request allocation information only when actually needed. Use DD statements to cause all volumes to be mounted when requesting allocation information.
- To obtain copies of individual VSAM data sets, use the EXPORT command.
- VSAM data volumes may be dumped and restored independent of the associated catalog(s) (BCSs).
- After restoring a VSAM data volume, BCS mismatches should be resolved and all data sets on the volume made current as required.

Procedures for Catalogs

This section covers the procedures associated with catalog functions explicitly, and with the catalog as a data set. An integrated catalog facility catalog can be split, merged, reorganized, exported, and imported.

All data sets cataloged in integrated catalog facility catalogs use access method services for catalog and scratch functions. Authorized programs such as IEHPROGM can be restricted to selected users.

The procedures for structuring the catalog using the DEFINE and ALTER commands must follow integrated catalog facility catalog requirements (see *Access Method Services Reference*).

The backup and recovery procedures for integrated catalog facility catalogs should be planned early in your installation or conversion process. (For complete information on backup and recovery, see *Catalog Administration Guide*.) Backup can be accomplished by EXPORT of the catalog or by a dump of the catalog volume. Because the BCS cannot be individually restored from a volume dump, EXPORT is usually preferred.

If recovery of a BCS is necessary, the DIAGNOSE command should be used to locate discrepancies between the BCS and VVDSs for the VSAM data sets. The DIAGNOSE command cannot be used for non-VSAM data sets, so the VTOCs of associated data volumes must be compared to the BCS (either manually or through a user program) to find the discrepancies. An alternative to this approach is to use SMF data to determine what catalog activity has taken place since the restored copy was created.

Discrepancies between the BCS and the data sets as they actually exist should be resolved using the RECATALOG parameter of the DEFINE command and the DEFINE NONVSAM and/or DELETE NOSCRATCH commands, as appropriate.

The development, testing, and documentation of the new backup and recovery procedures are a significant portion of the installation effort.

The catalog installation effort is determined chiefly by the number and types of catalogs in the installation, and the documentation of existing procedures.

Procedures for the physical conversion of existing catalogs must be developed for use during installation. Each catalog conversion must be scheduled and coordinated with related procedural changes. (See "Catalog Conversion with Name Retention" on page 70.)

Procedures for Catalogs Summary

With the replacement of VSAM catalogs and OS CVOLs with integrated catalog facility catalogs, many of the machine procedures for managing catalogs must be changed either to accommodate the new architecture or to use the new functions. The following items summarize the above discussion and provide additional guidelines.

- When converting catalogs containing generation data groups (GDGs) to an integrated catalog facility catalog on a different volume, model DSCBs must be re-created for the integrated catalog facility catalog volume.
- When converting VSAM catalogs, ensure that each owned data volume contains enough space for a VVDS and enough DSCBs to contain converted, suballocated VSAM components.
- To activate data set protection, be sure to protect the new catalog used in the conversion.
- When converting catalogs containing RACF-protected VSAM data sets, the RACF profile should consist of a component name and its catalog's volume serial number.
- Each catalog should have a documented, tested backup and recovery procedure.
- Integrated catalog facility catalogs (the BCSs) can be backed up by using EXPORT or REPRO. The frequency of backup is determined by the amount of catalog activity and the acceptable recovery time. An exported copy can be restored by IMPORT.

- The VVDS reflects the current status of VSAM data sets. The VVDS can be backed up and restored only as part of a total volume operation, not by using EXPORT or REPRO.
- As with the VTOC, if the VVDS is unusable, it can be rebuilt by recovering (possibly through IMPORT) each of its entries.
- After either a BCS recovery or a data volume recovery, the BCS should be compared with the associated data volumes. Use the DIAGNOSE command for VSAM data sets. For non-VSAM data sets, the installation must provide its own program or procedure to compare the BCS.
- Any discrepancies that are discovered should be resolved using the RECATALOG parameter of the DEFINE command and the DEFINE NONVSAM and/or DELETE NOSCRATCH commands.
- To reorganize the BCS, EXPORT a copy, use the DELETE RECOVERY command, use DEFINE to redefine the BCS, and IMPORT the copy using the INTOEMPTY parameter. The DELETE RECOVERY and DEFINE steps are needed only because the IMPORT command alone would cause the control area size to change if the control area size were less than a cylinder and if the IMBED option were in effect.
- The contents of one integrated catalog facility catalog can be merged into another integrated catalog facility catalog using REPRO MERGE. Entries from the source catalog can be selected individually by name or by high-level name qualifier.

Security

Both VSAM passwords and RACF are supported for integrated catalog facility catalogs. Unless the security system is modified, existing user procedures should not need many changes.

If RACF is used to protect VSAM data sets, the name of each protected component should be unique within the system. (See *VSAM Administration Guide*.)

Because a VSAM data set is accessible only through its catalog, protection of the data set requires protection of its catalog entry, and therefore, protection of its catalog (the BCS).

With the generic profile checking facility of RACF Release 5, an installation can consolidate the access authorization requirements of several similarly named and similarly used data sets under a single generic profile definition. A **generic profile** is used to protect a single cluster or a group of clusters that require similar access authority. A data set protected by a generic profile is recognized by RACF as different from a data set protected by a discrete profile.

For clusters cataloged in an integrated catalog facility catalog, a RACF generic profile will be used to verify access to the entire cluster, or any of its components. Discrete profiles for the individual components may exist, but **only** the cluster's

profile (generic or discrete) will be used to protect the components in the cluster. (Note: Profiles defined by ADSP processing during a data set define operation will be cluster profiles only.)

Data sets protected with discrete profiles are flagged as "RACF indicated." If a data set protected by a discrete profile is moved to a system in which RACF is not installed, no users will be given authority to access the data set. However, if the data set is protected with a generic profile, it is not flagged as "RACF indicated"; therefore, access authority is determined by normal VSAM password protection.

RACF will always be called to verify a user's authority to access both VSAM and non-VSAM data sets in an integrated catalog facility catalog if the data sets are protected by RACF. If the data sets are not protected by RACF, the data set passwords will be used to verify access authority. The integrated catalog facility catalog does not have to be RACF protected in order for its data sets to be RACF protected.

For more information on discrete and generic profiles, see *RACF General Information Manual*.

The holder of RACF alter authority or the master-level password for a catalog can gain access to the VSAM data sets in that catalog. Certain recovery operations for the VVDS require RACF alter authority or the master password of the catalog through which the VVDS is referenced. Application data with very stringent security requirements may justify a separate catalog because the holder of RACF alter authority for the catalog can gain access to the VSAM data sets.

Each catalog should be protected: at the master-level to prevent the deletion of VSAM data sets, and at the update-level if users are to be restricted to certain catalogs. Access lists for the catalog at the appropriate levels must be established.

The designated catalog owner is responsible for the security of the catalog. The owner has RACF alter authority for the catalog and maintains the access list. Users assigned to a catalog only need update access. Because performance tuning and backup and recovery operations require RACF alter authority for the catalog, the owner is also responsible for these functions.

The master catalog(s) should contain only the required system data sets, user catalog connectors, and aliases. To ensure this, the master catalog should be protected from update with very limited update access.

Erasing Residual Data

Facilities exist to cause the data component of a cluster or alternate index to be made unreadable when the data set is deleted. These facilities should be considered for use when the data set contains confidential or sensitive information, to prevent access to such data when the direct access space is reallocated for other use.

The access method services and TSO commands ALTER, DEFINE, DELETE, EXPORT, and IMPORT provide an ERASE parameter to specify that information should be erased when the VSAM data set is deleted.

RACF commands also provide for specification of an ERASE attribute in new or updated data set profiles (either generic or discrete). When VSAM data sets are deleted, the data component will be erased if either the integrated catalog facility catalog entry or the RACF profile contains the ERASE attribute.

The NOERASE parameter on an access method services or TSO DELETE command will not override the ERASE attribute specified in a RACF profile. For information regarding specification and use of this facility, see *RACF General Information Manual* and associated RACF publications.

System Resources

The system resources for supporting the catalog facilities include the processor, real storage, virtual storage, and I/O resources. Of these, the most important are virtual storage and I/O resources.

The system resources required to support catalog facilities depend directly on the number of catalogs (the configuration), and on the construction of the individual catalogs. These in turn depend on the objectives for the catalog environment, mainly for availability, security, and performance.

Virtual Storage Constraint Relief

The amounts given below for each release represent changes from the previous release. Each new release includes and adds to the gains made by the previous release.

Version 2 Release 1

Most of the catalog management modules and control blocks have been moved from the common service area (CSA) and link pack area (LPA) to a new catalog address space (CAS). This change frees approximately 410K bytes of storage (plus 20K bytes per open integrated catalog facility catalog¹) below 16 megabytes virtual.

¹ This CSA/SQA savings of 20K bytes is calculated for an integrated catalog facility catalog with two strings, data CI size of 4K bytes, and index CI size of 1K bytes. The savings would be larger for each catalog defined with more than two strings or additional buffers.

Version 2 Release 2

Approximately 1K bytes of virtual storage below 16 megabytes has been freed. With the expanded block size capability, you may realize an additional increase for new or redefined data sets if VSAM buffers are defined or allowed to default below 16 megabytes virtual.

Version 2 Release 3

Some catalog management modules have been moved from the link pack area (LPA) to the extended link pack area (ELPA). Also, some VSAM, open/close/EOV, and DADSM support has been moved above 16 megabytes virtual. As a result of these changes, approximately 350K bytes of storage below 16 megabytes virtual have been freed.

I/O Resources

The use of I/O paths, DASD devices, and DASD space in support of integrated catalog facility catalogs should be less than for other types of catalogs. The use of these resources depends chiefly on the control interval (CI) size, the number of buffers, and the share options specified for the integrated catalog facility catalogs.

A detailed understanding of the I/O resource use requires a knowledge of the integrated catalog facility catalog architecture and of VSAM as an access method (see *Catalog Administration Guide* and *VSAM Administration Guide*).

DASD Space and Units

Integrated catalog facility catalogs use DASD space efficiently. Generally, the BCS of an integrated catalog facility catalog is about one-quarter the size of a VSAM catalog for the same number of entries. When a VVDS is required for VSAM entries, it is usually small. For the number of components for each track the VVDS contains, see "Setting Up the VVDS" on page 66. Depending on user-specified options and data set names, an integrated catalog facility catalog may be about the same size as an OS CVOL for an equivalent number of entries.

Conversion to integrated catalog facility catalogs may not only free some DASD space because of better space utilization, but their performance characteristics may also allow catalogs to be consolidated onto fewer volumes. This gives an installation more flexibility in the use of space available on current catalog volumes, and may result in more usable space.

Defining the Installation Tasks

The organization of the tasks depends on the scope of the changes to be introduced. The following is a suggested order for the main installation tasks:

1. Realign catalog management and support responsibilities, if required.
2. Revise administrative procedures reflecting any realignment.
3. Develop, test, and document catalog performance guidelines.
4. Develop, test, and document catalog backup and recovery guidelines.
5. Develop, test, and document catalog conversion guidelines.
6. Analyze resource requirements for the integrated catalog facility catalog environment.
7. Design the catalog configuration.
8. Design individual catalog construction.
9. Revise and test VSAM data set procedures.
10. Revise and test catalog procedures to:
 - Define and alter catalogs using models
 - Reorganize catalogs
 - Back up catalogs
 - Recover catalogs and associated volumes or data sets
 - Convert VSAM catalogs and OS CVOLs to integrated catalog facility catalogs
 - Split and merge catalogs
 - List catalog entries
 - Explicitly create non-VSAM entries
 - Explicitly create aliases for user catalogs

Scheduling Catalog Conversions

Catalogs with high activity are candidates for early conversion (after the first production integrated catalog facility catalog is converted). Such catalogs typically support TSO functions.

A master catalog that contains only the required entries should be static. Such a master catalog may be a candidate for early conversion, to relieve the volume ownership of the system packs.

Recovery procedures during conversion of catalogs that contain VSAM data sets require an additional set of commands. Therefore, more planning and testing are needed, particularly when such catalogs support data base applications. It may take somewhat longer to replace these catalogs with production integrated catalog facility catalogs.

All systems that need access to the integrated catalog facility catalogs must have already installed the integrated catalog facility catalog support. It may be necessary to propagate such a system throughout the complex before an integrated catalog facility catalog is used in production.

The schedule for the actual conversion of existing catalogs should specify dates:

- By installation or location
- By system or complex (for example, all systems sharing the catalog)
- By application or set of applications

Remember that certain data set and catalog procedures must be implemented when a particular catalog or set of catalogs is converted.

The First Production Integrated Catalog Facility Catalog

Initially, the integrated catalog facility catalog-supported system may go into production with no production integrated catalog facility catalogs. This is particularly true in multiple processor installations because all systems requiring access to integrated catalog facility catalogs must have integrated catalog facility catalog support installed. When all processors in the complex have integrated catalog facility catalog support, conversion of production catalogs may start.

The first catalog to be converted to an integrated catalog facility catalog should be one that does not have stringent availability requirements. On the other hand, it should be one that has a considerable amount of activity so that meaningful production experience can be obtained. A TSO CVOL or similar catalog may be a good choice for the initial conversion. The new integrated catalog facility catalog can be frequently backed up and, if necessary, restored with minimal effort.

Post-Installation Assessment

After the conversion to integrated catalog facility catalogs is complete, the catalog environment should be reviewed to determine whether the objectives were met. Some refinements may be needed for the catalog configuration, construction, and procedures.



Chapter 7. Integrated Catalog Facility Catalog Conversion

Converting a VSAM catalog or OS CVOL to an integrated catalog facility catalog requires several steps:

1. Choose a conversion technique.
2. Check the CNVTCAT prerequisites.
3. Prepare for CNVTCAT processing.
4. Back up VSAM data sets, data spaces, and the catalog.
5. Set up the VVDS.
6. Set up the BCS.
7. Convert the VSAM catalog or OS CVOL.
8. Verify that the conversion has been successful.
9. Revise access method services procedures.
10. Examine application requirements.

Planning the Conversion of a VSAM User Catalog

The following are a few approaches that may be considered when planning the conversion from a VSAM catalog to an integrated catalog facility catalog.

Converting One Catalog at a Time

The simplest type of conversion is to convert one catalog at a time, using the access method services CNVTCAT command. In order to use the CNVTCAT command, certain prerequisites must be met. In some cases it is necessary to do some preparatory work to meet these prerequisites. The entire VSAM catalog and all its owned volumes are converted by executing the CNVTCAT command. Entries are built in the integrated catalog facility catalog to reflect the existing VSAM and non-VSAM data set entries. Existing data sets are not moved or renamed as part of the CNVTCAT process. (See "Procedures for Catalogs" on page 54.)

Converting One Volume at a Time

If each VSAM catalog controls only the volume it resides on, use the procedure described above.

If the VSAM catalog is recoverable, to perform volume conversion, use EXPORTRA ALL followed by IMPORTRA into the target integrated catalog facility catalog.

If a nonrecoverable VSAM catalog controls more than one volume, conversion on a volume-by-volume basis can be achieved only by explicitly exporting each data set from the VSAM catalog and importing them into the target integrated catalog facility catalog.

Converting on an Application, User, or Data Set Basis

These approaches use an explicit EXPORT and IMPORT sequence. (See "Procedures for Data Sets" on page 52.)

Because CNVTCAT provides a simple method of conversion, it is recommended for use whenever possible. With CNVTCAT, each VSAM catalog entry is converted to a corresponding integrated catalog facility catalog entry. Subsequent reorganization of the catalog structure can be achieved using the facilities of REPRO MERGECAT after experience with the integrated catalog facility catalog is acquired.

Checking the CNVTCAT Prerequisites

The prerequisites for the CNVTCAT technique are as follows:

- There must be sufficient space for a VVDS on each of the volumes owned by the converted catalog.
- There must be space for the BCS.
- There must be sufficient free space (especially if there are many suballocated VSAM data sets) in the VTOC of each volume owned by the converted catalog to allow for the new DSCBs that will be required.

The amount of space required for a VVDS is only a few tracks. If there is insufficient free space on the volume, it may be possible to release space for a VVDS by moving one or more VSAM or non-VSAM data sets to another volume.

The BCS can be placed on any volume, including:

- The volume containing the catalog to be converted
- Another volume that is owned by the catalog to be converted
- Any other volume

The space required for the BCS is usually about 25% of the space required for a similar VSAM catalog. However, both catalogs must exist simultaneously if CNVTCAT is to be used.

Each suballocated VSAM data set requires additional Format-1 (and possibly Format-3) DSCBs in the VTOC for each of its components following conversion. The BCS and each VVDS also require DSCBs. The total requirement is reduced by the number of existing DSCBs used for the VSAM data spaces that are released during conversion.

If the above prerequisites cannot be met, a certain amount of preparatory work is required before CNVTCAT can be used. In some cases, the amount of processing required may make the CNVTCAT approach undesirable.

Preparing for CNVTCAT Processing

To meet the prerequisite requirements for CNVTCAT processing, certain conditions must be established. The methods used to establish these conditions depend on the environment. A few of the possible combinations follow.

Full Volumes

When a catalog owns a VSAM volume with no free space for a VVDS and no opportunity to move non-VSAM data sets to another volume, EXPORT and IMPORT can be used to move VSAM data sets. Because the format of the portable data set is unchanged in the integrated catalog facility catalog environment, a data set exported from a VSAM catalog can be imported into an integrated catalog facility catalog.

Full Volumes with Some UNIQUE Data Sets

When a VSAM volume includes one or more data sets defined with the UNIQUE parameter, export enough data sets with the PERMANENT option to release space for a VVDS. The rest of the volume can be converted using the CNVTCAT command, and the exported data set can be imported into the target catalog. In some cases, the amount of space required for the VVDS may prevent the exported data set from being reimported to the original volume. In such cases, there is no alternative but to import the data set to another volume.

Full Volumes, No UNIQUE Data Sets, with Recoverable VSAM Catalog

If there are no UNIQUE VSAM data sets on a volume, all the data sets must be exported before any space can be released.

The entire contents of a volume owned by a VSAM recoverable catalog can be exported using EXPORTRA ALL. Although the integrated catalog facility catalog does not have a catalog recovery area and cannot be created with the RECOVERABLE attribute, it is still possible to use the IMPORTRA command to import the contents of a volume into an integrated catalog facility catalog. The volume's space may be reclaimed after the data sets are exported. This should release enough space for a VVDS, because even if there was no free space in the VSAM data space, a minimum of 1 cylinder was allowed for the catalog recovery area, which is not needed for an integrated catalog facility catalog.

Full Volumes, No UNIQUE Data Sets, with Nonrecoverable VSAM Catalog

If a volume is owned by a nonrecoverable VSAM catalog, export each of the data sets individually onto the volume and import them into the integrated catalog facility catalog. There is a possibility that all the exported data sets may not fit on the converted volume because of the space requirement of the VVDS. It may be possible to use CNVTCAT for some of the owned volumes to avoid the need for preparation of a lengthy EXPORT/IMPORT job stream.

Full VTOC

If the VTOC on a volume has insufficient free space to accommodate the extra DSCBs, it is necessary to rebuild the volume with a larger VTOC. VSAM data sets on such a volume should be exported with the PERMANENT option. Non-VSAM data sets should be dumped, using functions provided by Data Facility Data Set Services (DFDSS) or a similar function. If there is a VSAM catalog on the volume that owns other volumes, it should be unloaded using REPRO. A larger VTOC on the volume can then be built using functions provided by Device Support Facilities and the non-VSAM data sets restored. If a VSAM catalog was unloaded, it should be reloaded using REPRO and converted with CNVTCAT. Otherwise, an integrated catalog facility catalog should be defined. The exported data sets should then be imported into the integrated catalog facility catalog to complete the conversion.

Backing Up VSAM Data Sets, Data Spaces, Catalogs, and OS CVOLs

The CNVTCAT command deletes all VSAM data spaces and makes the source catalog unusable. As the conversion proceeds, integrated catalog facility catalog entries are built and DSCBs are created for the data sets that were suballocated. The conversion passes through a vulnerable period when data sets are not yet represented in the VTOC but their data spaces have been scratched. CNVTCAT ensures it has exclusive control of the VTOC during this process so there is no risk of assigning deleted space to another requester. However, if CNVTCAT fails during this vulnerable phase, the data sets are lost and conversion cannot be restarted. Similar considerations apply to the volume containing the VSAM catalog to be converted. After conversion has started for the catalog volume, the VSAM catalog is no longer accessible, because its VTOC DSCB is removed.

These facts make it necessary to back up all the VSAM data sets, data spaces, and the catalog itself prior to starting the conversion. Because conversion usually takes place when normal production systems are inactive, full dumps, using functions such as those provided by DFDSS, are the most convenient form of backup prior to conversion. The alternative backup procedure is to take an EXPORT copy of the VSAM data sets and unload the VSAM catalog, using REPRO. In either case, it is necessary to restart conversion by restoring all the VSAM entities on partially converted volumes.

OS CVOLs are not deleted by the conversion process. No special backup considerations apply to the OS CVOL, because the conversion process can be restarted if it is interrupted. For more information about converting OS CVOLs, see *Catalog Administration Guide*.

Setting Up the VVDS

The VVDS is an entry-sequenced data set (ESDS) with fixed characteristics. The only parameters that may be changed from their default values are the primary and secondary space allocation quantities. The control interval size is fixed at 4K bytes.

The VVDS contains information about the VSAM data sets on its volume. The number of records in the VVDS depends on the number of VSAM data sets and their type. The size of a record in the VVDS depends on the type of entry described and the length of the component name.

The first record in the VVDS is the VSAM volume control record (VVCR), which points back to all the catalogs that have data sets on the volume. The second record, the VSAM volume record (VVR), is the VVDS self-describing cluster record. For a guide to sizing the VVDS, see *Catalog Administration Guide*. The VVDS can go into secondary extents if it becomes full; however, to avoid fragmentation of a volume, choose an adequate primary extent size. If the VVDS is not explicitly defined, the default allocation of TRACKS (3 2) is used. This default allows for the following approximate numbers of components to be defined on a volume before filling a primary VVDS extent of 3 tracks:

IBM Device Type	Number of Components
3330 Disk Storage	70
3340 Direct Access Storage	40
3350 Direct Access Storage	100
3375 Direct Access Storage	220
3380 Direct Access Storage	280

To get the approximate number of VSAM data sets this represents, allow:

- 1 component for each entry-sequenced data set or relative record data set
- 2 components for each key-sequenced data set without IMBED
- 3 components for each key-sequenced data set with IMBED

Setting Up the BCS

The following items need consideration when setting up a BCS:

- Size of the BCS
- Number and placement of BCSs
- Tuning parameters for the BCS
- Naming convention for the BCS

The BCS contains records that describe the cataloged entities. The records are made up of smaller units known as subrecords, components, and cells. The cell is the smallest block of information.

To estimate the size of a BCS, it is useful to know the record sizes for different types of catalog entries. Actual record sizes depend on the:

- Length of the data set component names
- Number of volumes per data set
- Number of relationships between components
- Presence or absence of security information
- Number of alternate indexes
- Number of paths

The table below lists the approximate *minimum* lengths (in bytes) for various records, and the variables that affect the length of each record.

Record	Approx. Minimum Length	Variables Affecting Length
ALIAS	60	Length of association cell.
GDS	90	Number and length of volume cell(s); presence of association cell.
Non-VSAM data set	90	Number and length of volume cell(s); presence of association cell.
User Catalog Connector	90	Number and length of volume cell(s); presence of association cell.
GDG	260	Length of GAT cell; number and length of association and volume cell(s).
ESDS/LDS	230	Number and length of association and volume cell(s); presence of security cell(s).
RRDS	230	Number and length of association and volume cell(s); presence of security cell(s).
KSDS	400	Number and length of alternate index subrecords and volume cells; presence of association, relation, and security cell(s).
Alternate Index	170	Length of association, data name, index name, and volume cell(s).
Truename	70	Length of association cell.
Path	80	Length of association cell; presence of security cell.

The space parameters specified to provide for the estimated size of the BCS should take into account the:

- Data control interval size
- Device type
- Data control area size
- Embedded index (by default)

The relationship between these factors and the space requirements is the same as for a standard key-sequenced data set. Further details on space requirements for a standard key-sequenced data set will be found in *VSAM Administration Guide*.

Because most of the processing is random and no benefit is gained from large control interval sizes, the control interval size for the data component should be specified as 1K byte. Most data records are a few hundred bytes long, and a 1K-byte control interval size provides a useful compromise between minimizing data transfer time and reducing the incidence of a record spanning a control interval.

The control area size for the data component is effectively controlled by the secondary allocation quantity. If the secondary allocation quantity is:

- One track or less, the control area size is 1 track

- Greater than 1 track but less than 1 cylinder, the control area size is the same as the secondary allocation quantity
- One cylinder or greater, the control area size is 1 cylinder

The control area size should be sufficient to contain a maximum length record; the default maximum record length for the BCS (a spanned data set) is 32400 bytes. The following shows the minimum data control area sizes that are sufficient to contain a default maximum-length spanned record for a 1K-byte data control interval size:

IBM Device Type	Number of Tracks
3330 Disk Storage	3
3340 Direct Access Storage	6
3350 Direct Access Storage	3
3375 Direct Access Storage	2
3380 Direct Access Storage	2

If a smaller control area size is selected, the maximum record size must be reduced to fit in a control area by specifying the RECORDSIZE parameter with the DEFINE USERCATALOG command.

Because the BCS is a key-sequenced data set with the IMBED attribute, an additional track per control area is required for the sequence set records.

The process for estimating the space allocation for the BCS is to:

1. Estimate the total record size requirements in number of 1K-byte data control intervals.
2. Increase this quantity by the amount of free space required (for example, 20%).
3. Select an index control interval size and data control interval size based on the total data requirement and the device type.
4. Calculate the number of data control areas required.
5. Calculate total number of tracks (primary) required based on selected control area size, plus 1 additional track per control area for the sequence set records.
6. Specify the space allocation for the data component as TRACKS (primary secondary), where primary was calculated above and secondary is the selected data control area size.
7. Specify the space allocation for the index component as TRACKS (1 1).

Note: The above procedure assumes a reorganized BCS without unused space caused by control interval and control area splits. An empty catalog that was just defined contains records for the catalog cluster, with its data and index components and also the VVDS. Subsequent additions to the catalog are insertions, and may cause a significant number of control interval and control area splits resulting in an under used BCS. The amount of space required during

the initial period when CNVTCAT or DEFINE is used to build entries in the BCS may exceed the estimated size, possibly by a factor of 4. Reorganizing the BCS by using EXPORT and IMPORT reduces the space requirements to a value nearer the one resulting from the above estimate.

The definition of the BCS can include the specification of other parameters for selecting performance characteristics. BUFFERSPACE is specified directly by using the STRNO parameter. Initially, use the default value of STRNO(2). This can be increased if frequent enqueue contention occurs for the resource SYSZRPLW.catname. The major name is SYSZRPLW and the minor name is catname.

Converting the Catalog

The CNVTCAT process requires exclusive control of the catalog to be converted. If the process fails, a volume restore operation will be required. It is recommended that the conversion take place in an environment where other jobs cannot access the volume; for instance, in a stand-alone environment.

Following CNVTCAT, the user catalog connector for the source catalog is unchanged and aliases for the source catalog are left in place. To complete the conversion, remove the user catalog connector for the source catalog, using EXPORT DISCONNECT. This also removes the aliases. Redefine the aliases for the target catalog.

With a large number of alias entries, as in some TSO installations in which each userid is an alias of a catalog, the generation of DEFINE ALIAS commands to rebuild these alias entries may be automated. Use an online editor to process a LISTCAT output listing of the catalog prior to conversion, or a simple program to perform a similar function.

Catalog Conversion with Name Retention

An alternative two-step method for catalog conversion will retain the catalog name and associated aliases. The source VSAM catalog or OS CVOL is converted to an interim integrated catalog facility catalog, perhaps on another volume, and the source catalog is deleted. A target catalog with the same name as the source catalog and on the same volume is defined. The contents of the interim catalog are transferred to the target catalog, using REPRO MERGECAT. The interim catalog, now empty, may be deleted. For an example of this process, see *Catalog Administration Guide*.

Prior to conversion, the target BCS contains:

- Its own sphere record
- The BCS data and index records
- The VVDS sphere records

As each object in the source catalog is converted, new records are inserted into the target BCS. Apart from true name records with a generated name (that is, a name prefixed with VSAMDSET), these insertions take place in approximately

ascending collating sequence. The BCS may have a number of control interval and control area splits. Because the splitting process leaves both the control intervals and control areas half full, it is possible to produce a BCS that is only one-quarter full.

Listing the cluster entry for the BCS following conversion shows if control interval or control area splits occurred. To improve BCS space utilization and remove the fragmentation caused by control interval and control area splits, the BCS is usually reorganized following conversion. This is done by exporting the BCS and then importing it. The processing done by IMPORT results in the (basic catalog structure) requested FREESPACE parameter amounts being preserved.

Verifying the Conversion

The conversion can be checked in the following ways:

- Inspecting the output from CNVTCAT

CNVTCAT issues messages to indicate the start and successful completion of conversion for each entry in the source catalog. It also issues a message as each volume's VSAM data spaces are about to be deleted.

- Inspecting the output from EXPORT/IMPORT

If IMPORT is used for all or part of the conversion process, it gives the date and time the exported copy was made. Verify this against the EXPORT listing to ensure that an earlier copy of the data was not imported.

- Listing the VTOC of the converted volume

LISTVTOC FORMAT output from IEHLIST shows the successful conversion is completed as follows:

- The source catalog's data space entry is scratched.
- The data and index components of data sets that were formerly suballocated are in the VTOC.
- Converted data set entries have OPTCD = X'80'. The integrated catalog facility catalog entries have OPTCD = X'C0'.

The VSAM-ownership bit in the Format-4 DSCB is turned off. This can be checked by running a LISTVTOC DUMP or an AMASPZAP VERIFY on the DSCB.

- Listing the target catalog

The target catalog shows all the data set entries that were in the source catalog. Owned volumes have a VVDS entry instead of a volume record. The catalog cluster name always shows as 44 characters of zeros, reflecting the catalog's self-describing record key of binary zero. The catalog cluster record always appears first in the catalog listing.

- Running DIAGNOSE on the BCS and its connected VVDSs

DIAGNOSE allows the BCS and each of its VVDSs to be checked against each other for consistency. No entries should be flagged as in error.

If the LIST option is selected, the name and type of each entry in the BCS and VVDS are listed.

Revising Procedures for Integrated Catalog Facility Catalogs

The installation's backup and recovery procedures may handle the integrated catalog facility catalog environment as soon as it becomes part of the production system. Users of the VSAM recoverable catalog may find that their existing procedures no longer work. A new backup and recovery procedure should be designed and tested before using the integrated catalog facility catalog.

Production job streams may contain the names of catalogs in STEPCAT or JOBCAT DD statements. If any catalogs have changed their names, these job streams must be revised. (See "Converting the Catalog" on page 70.)

Production job streams or TSO CLISTs may contain access method services commands. Most of these commands function normally without alteration. However, the DEFINE SPACE, EXPORTRA, LISTCRA, and RESETCAT commands perform differently or do not apply to the integrated catalog facility catalog environment. For further information about these access method services commands, see *Access Method Services Reference*. If catalog names are included in commands contained in production job streams, these may need to be revised.

VSAM UNIQUE data set suballocation of data space does not exist in an integrated catalog facility catalog. All objects will be defined UNIQUE. Therefore, the UNIQUE and SUBALLOCATION parameters of the DEFINE commands should not be used. If used, however, they are ignored unless specified along with a conflicting parameter (such as REUSE); in that case, a message is issued saying that the parameters are inconsistent.

Space requests for unique allocations (for all integrated catalog facility catalog allocations) are not rounded to a cylinder as they are with VSAM catalogs. Thus, applications that have used more than the requested space in VSAM catalogs may fail for lack of space when the data sets are deleted and redefined for integrated catalog facility catalogs. If allocation is in tracks, they will be in contiguous tracks. Suballocation is not allowed for reusable data sets for integrated catalog facility catalogs. In addition, secondary extents of such data sets on the first volume are not released when they are opened for reuse. Secondary extents on any other volume than the first volume will be released.

The volume-oriented commands EXPORTRA and IMPORTRA (as well as LISTCRA) are not available with integrated catalog facility catalogs. If users have relied on these commands to back up their data sets, then the backup procedures must be changed.

If the installation has used VSAM volume ownership to restrict VSAM allocations by volume, these procedures must be changed.

JCL does not allocate VSAM data sets, and control of VSAM allocations (through programming) can be done in the DADSM exits.

VSAM data sets are treated as unique data sets under the integrated catalog facility and cannot be deleted while allocated. The FILE parameter should be coded in the DELETE command, along with the corresponding DD statement, with DISP = SHR.

For Installations with VSAM Catalogs: The procedures for creating catalogs must be changed, because the default definition is for VSAM catalogs; in addition, default parameters of the DEFINE command may not be appropriate for performance objectives. Catalogs that are used as models for other catalogs must also be redefined.

Examining Application Requirements

No application using standard VSAM macro calls should require modification following the conversion to integrated catalog facility catalogs.

Test the following types of programs in the integrated catalog facility catalog environment to determine whether they operate correctly:

- Programs that issue SVC26
- Programs that use CATLG and CAMLST

Programs that read catalogs as data sets need careful attention. If the functions these programs perform are still required in the integrated catalog facility catalog environment, the programs must be changed to allow for the new record formats. Because catalog information resides in both the BCS and the VVDS, such programs may require substantial rewriting. The access method services PRINT command can be used to display the contents of BCS and VVDS records to show typical layouts. (See *Access Method Services Reference*.)

Planning the Conversion of a VSAM Master Catalog

The considerations for converting a VSAM catalog to an integrated catalog facility catalog also apply to converting a VSAM master catalog. There is, however, an additional restriction: an active master catalog cannot be converted if it is currently in use by the system.

One approach to converting a VSAM master catalog is to use a temporary, or alternate, master catalog while the VSAM master catalog is being converted. A pointer to this alternate master catalog is put in a SYS1.NUCLEUS member, SYSCATnn. At IPL time, the operator specifies SYSCATnn, and the alternate catalog is used as the master catalog. Any existing SYSCATnn member may be used (including the default member SYSCATLG) or a new SYSCATnn member can be created. For further information on SYS1.NUCLEUS, see *System Generation*.

The temporary (alternate) master catalog must be built by the installation, and needs to contain only system data set entries and connectors to the source and target catalogs. The page and swap data sets defined in the temporary catalog can be just large enough to run the conversion, and should be independent of the catalog being converted.

When a fully operational alternate master catalog is generated for the system, new page space data sets for the alternate master catalog must be allocated because the old ones cannot be utilized. When the master and alternate master catalogs are functioning correctly, the VSAM catalog may be converted. For information about how to generate an alternate master catalog, see *Catalog Administration Guide*.

An alternative approach is to use a system that does not use the source catalog of the master catalog being converted; for example:

- An alternative processor that can access the volume that contains the catalog
- An operating system that can be restored to the processor that uses another alternate master catalog

Planning the Conversion of an OS CVOL

The following considerations pertain to OS CVOLs:

- The OS CVOL itself (that is, the SYSCTLG data set), which contains catalog entries for non-VSAM data sets and GDG data sets
- A non-VSAM entry in the master catalog for the OS CVOL, which normally has the name SYSCTLG.Vvolser
- One or more aliases for the OS CVOL in the master catalog, which direct searches for data sets with a high-level qualifier matching the alias name to use the OS CVOL

The conversion process includes constructing an integrated catalog facility catalog that is equivalent to an OS CVOL, and updating the alias entries to point to the integrated catalog facility catalog. Use the CNVTCAT command of access method services to build an integrated catalog facility catalog to contain all the entries of the OS CVOL. The source catalog is unchanged by the conversion, so there is no special requirement to back up the OS CVOL. The target catalog does not need to be empty and OS CVOLs can be merged during conversion. Otherwise, each OS CVOL can be converted to an integrated catalog facility catalog and the integrated catalog facility catalog tested for function and performance before doing any reorganization. The access method services REPRO command may then be used to split or merge parts of different integrated catalog facility catalogs.

No method of altering the alias entries in the master catalog to point to the integrated catalog facility catalog is provided. All aliases must be deleted (either individually or by deleting the SYSCTLG.Vvolser NVSAM entry, which takes

with it all the aliases), and new aliases must be defined for the integrated catalog facility catalog.

If the number of alias entries is large, as in some TSO installations in which each userid is an alias of a catalog, then the generation of DEFINE ALIAS commands to rebuild these alias entries may be automated. Use an online editor to process a LISTCAT output listing of the catalog prior to conversion, or use a simple program that performs a similar function.

Where an OS CVOL catalog contains entries for a GDG, a model DSCB for new GDG members normally resides on the same volume as the catalog. If the catalog volume is changed during conversion, care must be taken to identify and move these model DSCBs so they reside on the integrated catalog facility catalog.

Following conversion, revise procedures to allow for the integrated catalog facility catalog. Using the corresponding access method services functions, revise existing job streams that used IEHLIST, IEHPROGM, or IEHMOVE to manage the OS CVOL. For more information about converting OS CVOLs, see *Catalog Administration Guide*.

After the integrated catalog facility catalog is working correctly, the SYSCTLG data set and its master catalog entry can be deleted.



Chapter 8. Integrated Catalog Facility Catalog Backup and Recovery

Because most installations depend on the availability of cataloging facilities to run production job streams and to support online users, a catalog failure can cripple an installation's capability to do useful work. Therefore, backup and recovery strategies are crucial. An installation's recovery methods must be suited to the availability requirements of its various kinds of data, and its backup methods must support the chosen recovery methods.

The ability to respond appropriately in a recovery situation requires advance preparations. Your installation should test and document backup and recovery procedures for all catalogs.

Recovering a catalog (that is, the basic catalog structure) consists of making one or more data sets *available* for application processing, at an *acceptable level of currency*, and within a *reasonable amount of time*. Catalog recovery is usually needed because of either a damaged index component of the catalog or a hardware failure.

Creating a catalog backup consists of copying the catalog in a format suitable for recovery. Backups provide protection against hardware, software, and human errors, against disasters, and against deliberate destruction. Backups can also be used as archives to meet legal requirements and unanticipated business needs.

Data Error Detection and Diagnosis

In an actual recovery situation, you must first recognize the need to recover data and understand the source and impact of the data error. Only after this type of analysis can you choose appropriate recovery actions.

Structured analysis of the messages and information accompanying a failed attempt at data access usually can quickly identify the system function that failed and the data set or control structure most likely in error. This initial error detection should determine whether a data error caused the failure and whether the most likely failing structure is the VTOC index, VTOC, VVDS, catalog, other system data set, or user data set. You can then execute the diagnostic routine appropriate to the failing structure and find out whether there is an error and, if so, how much damage there is. If the suspected failing structure is a catalog, you can use any of five access method services commands for this diagnosis: DIAGNOSE ICFCATALOG, EXAMINE, LISTCAT VOL, LISTCAT ALL, or PRINT. If, however, damage to the catalog prevents these commands from working, you may need to use another diagnostic tool such as Data Facility Data

Set Services PRINT. A data error that is due to a hardware I/O error can also be diagnosed from the associated messages and from LOGREC information.

When only a few entries are in error, some people think that they can simply repair the catalog. However, trying to fix parts of a damaged catalog rather than replacing it with a backup is often impossible and usually only makes the problem worse. The recommended response to a damaged catalog is full catalog recovery, which includes updating a backup copy so that it contains information about data sets up to the point of failure.

Full Catalog Recovery

How can you update an old copy of a catalog so that it contains information about data sets up to the point of failure? You need to use various sources of information and various techniques, depending on the information source.

Sources of Current Information about Data Sets

The catalog is an inventory of permanent data sets stored on DASD or on tape. It contains the volume serial number(s), the device type (or an indirect pointer to a device type), and usually the location where the data set is described (TTR for the DSCB or the file number on tape). You may have several sources of current information about data sets:

- Volume table of contents (VTOC) and labels

The definitive source of information about a cataloged data set is on the volume that contains the data set: the VTOC for DASD and the data set label for tape. For permanently mounted DASD, this source of information is readily available. For tape, however, and to a lesser degree for mountable DASD, the data volumes are not a satisfactory source of information. The interactive storage management facility (ISMF) can provide information about cataloged data sets.

- Data and storage management applications

If your installation is using a data management or storage management system (such as Data Facility Hierarchical Storage Manager), there is probably an inventory of data sets that have been actively managed, or of data sets associated with a particular device type or category. Although not complete, these inventories may have data that is valuable for catalog recovery. Of course, the inventory will need to be reconciled with the entries in the recovered catalog; the DFHSM commands HLIST and AUDIT can be used for this purpose. For information about these commands, see *DFHSM System Programmer's Guide*.

- System management facilities (SMF)

Catalog management logs all changes to the SMF data set. If you record and save the SMF record types for catalog changes, you will have the information

you need for updating a recovered catalog. Your installation must develop its own programs and procedures to use this source.

Full Catalog Recovery Procedure

The catalog recovery procedure consists of preparing to recover from a damaged catalog, updating the recovered catalog so that the information reflects the data sets up to the point of failure, and completing the recovery.

Preparing for the Recovery

1. As of MVS/XA Data Facility Product Version 2 Release 3, aliases of an integrated catalog facility catalog are automatically written to the portable data set during EXPORT processing; these aliases can be defined or listed when the catalog is subsequently imported. On other systems, however, you need to list the aliases of the damaged catalog from each of the master catalogs (in a multisystem environment) to which it is connected; these aliases may have to be reestablished later.
2. Deny access to the damaged catalog from all systems. This can be done by locking the catalog (see "Enhanced Integrated Catalog Facility Catalog Recovery" on page 80). Systems that do not have this ability need to quiesce or terminate user sessions to prevent users from regaining access until the recovery has been completed.
3. Dump the SMF data sets on all systems that had access to the catalog. These records may be needed for updating the recovered catalog if user programs are available to utilize the data.
4. Save a copy of the damaged catalog for future diagnostics. Use the DFDSS logical data set DUMP command.
5. Save the contents of the catalog in a readable format. Use the DFDSS command PRINT by tracks to save both the data and the index components.
6. Save the VVDS for the catalog. Use the DFDSS command PRINT by tracks.
7. Delete the damaged catalog from the system by using the access method services command DELETE RECOVERY. This command also deletes the aliases.
8. Restore the catalog from your backup by one of the following methods:
 - Use the access method services IMPORT LOCK command. If you also specify the ALIAS parameter with the IMPORT command, the aliases that were exported can be defined or listed.
 - Use the DFDSS Version 2 Release 3 logical data set RESTORE command. Aliases are automatically restored when DFDSS is used. For details, see *DFDSS: User's Guide and Reference*.

Updating the Recovered Catalog

This section of the recovery procedure assumes that you have chosen to update the new catalog by using the VTOCs and other such sources of information. The approach is first to create catalog entries for uncataloged data sets and then to remove obsolete catalog entries.

1. Determine which volumes contain data sets that should be represented in the recovered catalog.
2. List the data sets on each volume. This can perhaps be done most easily by using the interactive storage management facility.
3. Determine whether each data set in a list is cataloged.
4. Catalog each data set that is uncataloged. Use the the access method services DEFINE RECATALOG for VSAM data sets and DEFINE NONVSAM for non-VSAM data sets.
5. List the cataloged data sets. This can perhaps be done most easily by using the interactive storage management facility.
6. Identify the nonexistent data sets.
7. Delete the obsolete catalog entries. Use the access method services DELETE VVR for VSAM data sets and DELETE NOSCRATCH for non-VSAM data sets.

Completing the Recovery

1. Check the status of the catalog by using the access method services command DIAGNOSE.
2. Rebuild the aliases for the catalog if they were not defined during import processing.
3. Back up the catalog to start a new recovery cycle.
4. Restore access to the catalog from all systems by unlocking the catalog.

Enhanced Integrated Catalog Facility Catalog Recovery

A catalog that is being recovered *must not* be accessed during the recovery. Before enhanced integrated catalog facility catalog recovery (available with MVS/XA Data Facility Product Version 2 Release 3), if you needed to repair or recover a catalog, you had to deny access to it by quiescing the system or terminating user sessions that were oriented to it. After the recovery you had to restart processing on the catalog.

Now you can deny access to an integrated catalog facility catalog during recovery with minimal impact on the surrounding subsystems. In particular, you do not have to *terminate* subsystems like IMS/VS and CICS/OS/VS that are oriented to

an integrated catalog facility catalog when you need to recover or repair the catalog. Rather, the catalog locking facility of enhanced integrated catalog facility catalog recovery guarantees the integrity of the catalog by causing requests to the locked catalog to fail. Furthermore, enhanced integrated catalog facility catalog recovery *automatically reorients the subsystems or users to the catalog after recovery.*

Catalog Locking

Within one system, enhanced integrated catalog facility catalog recovery makes it possible to lock integrated catalog facility catalogs during recovery. Catalog locking restricts access to the catalog so that authorized recovery personnel can perform the recovery while the catalog remains inaccessible to other users. Attempts to access the catalog return a code indicating that the catalog is temporarily unavailable. After recovery has been completed, recovery personnel can unlock the catalog so that normal operations can resume.

You can lock and unlock a catalog by using two parameters, LOCK and UNLOCK, with the access method services commands ALTER, DEFINE USERCATALOG, and IMPORT. These parameters apply only when the target catalog is an integrated catalog facility catalog. You should revise your existing procedures for denying access to the catalog by adding steps to lock and unlock the catalog during recovery. You need not change your procedures for catalog recovery, however; you can still use the access method services command DELETE RECOVERY to delete a damaged catalog before importing a backup copy.

ALTER LOCK locks an existing integrated catalog facility catalog; ALTER UNLOCK, as the last step of catalog recovery, unlocks the catalog. Access is allowed as soon as UNLOCK has been processed. DEFINE USERCATALOG LOCK or DEFINE USERCATALOG UNLOCK controls the setting of the lock attribute for a newly defined integrated catalog facility catalog. After the DELETE RECOVERY of an integrated catalog facility catalog, you should define the new catalog with the LOCK parameter to prevent users or online systems from using the catalog during recovery; under most other circumstances, UNLOCK is appropriate. IMPORT LOCK and IMPORT UNLOCK set the lock attribute on or off for the integrated catalog facility catalog being imported. Unless LOCK is specified on the DEFINE or IMPORT commands, the catalog will be defined as unlocked and will be accessible to users as well as recovery personnel.

Locking Authorization

The IGG.CATLOCK facility, in conjunction with normal security checking, controls who may invoke the lock attribute and who may access the locked catalog. Only users who have authority to the IGG.CATLOCK facility can lock an existing unlocked catalog or a newly defined catalog, unlock an existing locked catalog, or access and repair a locked catalog. In addition to authority to the IGG.CATLOCK facility, normal RACF checking or password validation is required. If your installation does not use RACF or an equivalent product, you must supply a routine to handle calls to IGG.CATLOCK.

For RACF systems, establishing authority to the IGG.CATLOCK facility requires a RACF RDEFINE command for the profile IGG.CATLOCK of class type FACILITY; it also requires a PERMIT command with read access to that profile for one or more users. RACF systems before Release 1.7 must add the FACILITY class to the class descriptor table (CDT). Non-RACF systems (including a user-coded MVS router exit) must be capable of interpreting and processing the profile IGG.CATLOCK of class type FACILITY.

To ensure the integrity of integrated catalog facility catalogs, installations should review the number of users who have authority to the IGG.CATLOCK facility and restrict this access to as few individuals as possible. Because normal processing of catalogs does not require authority to the IGG.CATLOCK facility, most users can be limited to update access on their catalogs if they must be able to define and delete objects in the catalog.

System management facilities (SMF) logs all alterations of catalog records, including the lock attribute in the catalog's self-describing records. This logging provides an audit trail of users who perform catalog recovery.

Considerations in a Shared Environment

For catalogs shared between processors, automatic system reorientation to a rebuilt catalog occurs only on the processor(s) with enhanced integrated catalog facility catalog recovery.

On systems without enhanced integrated catalog facility catalog recovery, locked catalogs are accessible; therefore, unpredictable results may occur if a processor tries to access a catalog that has been rebuilt on the same volume by another processor. In this case, standard quiesce procedures should be followed; you must close the catalog by terminating all tasks oriented to it.

Some online systems cannot adequately recover from attempts to access data bases cataloged in a shared catalog that is locked for recovery purposes. These data bases will need to be taken offline. The goal is to prevent OPEN, CLOSE, EOVS, ALLOCATE, and DEALLOCATE from accessing the locked catalog while it is being recovered.

Restrictions

For a catalog to be a candidate for enhanced integrated catalog facility catalog recovery, no job may be allocated to the catalog with a disposition of OLD. This disposition of OLD could come from JOBCAT or STEPCAT statements or DD cards, or from dynamic allocation of the catalog. This dynamic allocation can result from the use of OUTDATASET with most access method services commands. Specifying a disposition of SHR on a DD card and then referring to the DD card with the OUTFILE keyword bypasses this dynamic allocation.

Recovering a Master Catalog as a User Catalog

You must have a procedure for recovering the master catalog, especially in a single-system installation. The master catalog is critical because it is required to IPL and run an MVS system.

It is a good practice to restrict the use of the master catalog to objects that must reside there: system data sets, user catalog connectors, and aliases. This practice ensures that recovery of the master catalog consists only of CONNECT and DISCONNECT operations and the definition and deletion of aliases, functions that can easily be done against an active master catalog.

The master catalog cannot be recovered while it is an active master catalog, only as a user catalog. There are two approaches to accessing a master catalog as a user catalog.

In a single-system installation, the recommended approach is to use an alternate master catalog. This copy may or may not be the same as the master catalog to be recovered, but it must have been previously built. It cannot have the same name as the master catalog to be recovered. This approach requires that you re-IPL and specify the alternate master catalog using the IPL options.

In a multisystem environment, the recommended approach to recovery is to use a system other than the one with the failing catalog. The volume of the failing catalog must be on a unit accessible to another system, and the failing catalog must be connected to the master catalog of the recovery system. Therefore, the master catalogs must have different names and must be on volumes that have different serial numbers.

If you are using catalog locking as part of recovering an integrated catalog facility master catalog that is shared with another processor, you must do so from the processor to which the catalog is connected as a user catalog. Using the QUIESCE operator command, you should quiesce the processor that uses the catalog as a master catalog during the recovery. You cannot recover master catalogs to a different volume. Also, you must unlock a master catalog before doing an IPL; if you do not, the IPL will fail.

Recovering Individual Catalog Entries

Sometimes diagnosis of a catalog access problem shows that certain entries are missing from the catalog or have been damaged. The catalog is structurally sound and there are no hardware or media failures. Only the logical content of some catalog records is affected.

Recovery, in this case, requires removing any damaged catalog records (using the access method services DELETE commands for specific entries) and introducing new records to represent the existing data sets correctly.

However, full catalog recovery must be used if the cluster record or the data or index component truname record for the catalog is missing or damaged.

Catalog Backup

There is no substitute for adequate catalog backup. The backup procedures you choose depend on the recovery procedures you have chosen and plan to use if the need arises.

Creating backups and keeping records for backups require considerable system or human resources or both. As desirable as it might be to have a current copy of every data set, this is seldom possible because of constraints on resources and time. It is very important, however, to be able to recover a catalog without having to recover all the data sets represented in the catalog.

The access method services EXPORT command can be used to create catalog backups. Unlike Data Facility Hierarchical Storage Manager and Data Facility Data Set Services, using IMPORT for catalog recovery allows you the flexibility of recovering to another volume and a different device type. DFHSM has the advantage, however, of *automatically* making backups whenever a catalog is changed.

Before you back up the catalogs that control access to critical data sets, you should follow the steps listed in the following section to ensure that the copy about to be made will be usable if it is ever needed for recovery.

Steps in the Backup Process

Whatever backup technique you use, the backup process follows the same logical sequence of steps during preparation, execution, and administration and record-keeping.

Preparation

1. Decide on a technique for backing up each catalog. (All catalogs should be backed up.)
2. Decide whether multiple backup copies are needed.
3. Select the medium onto which the backups will be dumped (either tape or DASD).
4. Select the names to be used for the backup catalogs.
5. Set up the backup jobs or alter parameters for automatic backup.
6. Schedule the backup jobs or alter the start time for automatic backup.

Execution

1. Release backup jobs (if necessary) to allow automatic backup to begin.
2. Validate and record the contents of the VTOC(s).
3. Validate the integrity of the catalog that is to be backed up.
4. Serialize the use of each catalog as it is backed up.
5. Copy each catalog.
6. Release serialization of the catalog after it has been backed up.

Administration and Record Keeping

1. Check for successful backup of each catalog.
2. Provide notification of unsuccessful backup attempts.
3. Indicate that the catalog was successfully backed up, if appropriate.
4. Associate the name of the catalog with its backup copy.
5. Record the location of the backup copy.
6. Check the backup copies for readability (if required).
7. Create a duplicate backup, if appropriate.
8. Release any backup copies no longer needed.
9. Inventory and store the backup copy, if the backup volume is portable.

For further information on catalog backup and recovery, see *Catalog Administration Guide*.



Part 3. Planning for the Indexed VTOC



Chapter 9. Indexed VTOC Overview

An indexed volume table of contents (VTOC) is a combination of a VTOC in the OS format and an index for direct access into that VTOC. The VTOC index contains free space information as well as pointers into the VTOC to improve performance. This improvement avoids lengthy sequential-keyed searches, which make the device and channel unavailable; it also manages free space information, thus reducing the number of I/O operations to obtain or release space.

The VTOC index is optional. It can be built for an existing VTOC and later "switched off," with or without deletion of the index. This allows DADSM requests to be processed through the index, or directly through the VTOC without using the index. For a description of the VTOC index, see *System—Data Administration*.

The VTOC is constructed of 96-byte unblocked records with hardware keys. If an allocation request is made for an existing cataloged data set, a pointer to that data set's Format-1 DSCB may exist in the catalog entry. (A DSCB is a data set control block label for a data set on direct access storage.) A request to read the Format-1 DSCB reads the DSCB directly, using the track record (TTR) in the catalog. In other cases, to find a Format-1 DSCB for a data set, the hardware may have to perform a sequential-keyed search of the VTOC. This type of access makes the device and channel unavailable for the duration of the search, with the length of time related to the number of DSCBs in the VTOC. As the size of volumes increases, the number of Format-1 DSCBs increases, aggravating the problem. To locate free space on the volume requires a scan of the Format-5 DSCBs which, when multiple-chained, may be scattered through the VTOC, resulting in several I/O operations to the VTOC to satisfy the allocation request.

A request to read a Format-1 DSCB from an indexed VTOC when a TTR is not available results in a search of the VTOC index. The DSCB is then read directly without searching the entire VTOC. Free space on the volume is obtained from the index without searching long chains of Format-5 DSCBs. The indexed VTOC minimizes the disadvantages of the old VTOC but is compatible.

Structure

This section describes the relationship between the VTOC and its index and the internal structure of the index itself. Further information may be found in *System—Data Administration* and *DADSM Diagnosis Reference*.

Relationship between the VTOC and Its Index

The VTOC index is a specialized data set that resides on the same volume as the VTOC it refers to. It has a Format-1 DSCB in the VTOC that contains its name and extent information. The index must adhere to the naming convention 'SYS1.VTOCIX.xxxxxxxx', where 'xxxxxxx' can be anything valid in a data set name and is generally the volume serial of the volume containing the VTOC and its index. To avoid ENQ contention, it must be unique within the system.

The relationship of a VTOC to its index is shown in Figure 7. For a complete description of the parts of an index, see *System - Data Administration*.




Figure 7. Relationship of a VTOC to Its Index

The VTOC of an indexed VTOC volume also has several special characteristics:

- The Format-4 DSCB has its DOS bit “on” to assist compatibility and portability between indexed and nonindexed VTOCs, and between systems with and without the software support for indexed VTOCs.
- The Format-4 DSCB also has a bit called the index bit “on” to indicate that the volume contains an active VTOC index to be used with the VTOC.
- The number of Format-0 DSCBs is no longer necessary and is not maintained in the Format-4 DSCB.
- The pointer in the Format-4 DSCB to the highest Format-1 DSCB is set to point to the last DSCB slot in the VTOC extent that describes VTOC space.
- One Format-5 DSCB exists and is empty to provide a basis for conversion back to a nonindexed format, a process which requires reconstruction of Format-5 DSCB information.

As an example of how the two structures relate to each other, it is possible to access the VTOC if the index becomes unusable. However, it is not possible to access data sets on the volume if the VTOC index is accessible and the VTOC is unusable. The index is merely a means of accessing the VTOC. The VTOC must be read if the data set is to be used by standard access methods.

VTOC Index Structure

Internally, the VTOC index looks very much like a VSAM index, but it is actually physical sequential in organization. It contains two types of records:

- Space maps describing space in the index, in the VTOC, and on the volume
- VTOC index entry records containing entries for data sets and pointing to Format-1 and -4 DSCBs in the VTOC

Common VTOC Access Facility (CVAF)

CVAF provides a set of macros for accessing indexed and nonindexed VTOCs from user programs and performing some system functions. CVAF allows the user:

- To access the VTOC directly or sequentially
- To access the space maps in the VTOC index
- To test whether the system has the indexed VTOC feature installed, or a particular volume has an indexed VTOC on it

Other macros support initialization of indexed VTOC processing on a volume and recording of recovery data for system use.

Details of the syntax of the CVAF macros, return codes, CVAF internals, and diagnosis guidance are delineated in *System – Data Administration*, and *CVAF Diagnosis Reference*.

If you use CVAF, you are responsible for performing resource serialization on the VTOC and VTOC index, because CVAF assumes that the caller does this. The normal VTOC RESERVE name should be used (the major name is SYSVTOC, the minor name is volser). Programs that use RESERVE with this name must be authorized by the authorized program facility (APF) and reside in an authorized library. Read-only access to the VTOC or index should be prefixed with a shared RESERVE. Update access requires an exclusive RESERVE. The resource should be held for the minimum length of time. System performance may be adversely affected if the VTOC is unavailable for a long time (particularly in a shared DASD environment, where the entire volume is made unavailable by a RESERVE). A shared RESERVE holds up other processing if that processing requires updates to the VTOC or index, or if the volume is accessible from another processor.

A discussion of conversion of VTOC-reading user programs to use CVAF is included in “Planning for User Programs That Read VTOCs” on page 93.



Chapter 10. Indexed VTOC Conversion

Migration procedures to indexed VTOCs are very flexible. A system can contain only indexed VTOC volumes, only nonindexed VTOC volumes, or a mixture of both. This allows a gradual migration for the volumes an installation wants to convert to the indexed VTOC format.

A VTOC with an index can be created by Device Support Facilities when a volume is initialized with the Device Support Facilities INIT command, or an existing volume containing a VTOC may have an index built by the Device Support Facilities BUILDIX command. Indexed VTOCs may reside on any direct access storage device (DASD) supported by the host system control program. For further information, see *Device Support Facilities User's Guide and Reference*.

Another important consideration is the status of user programs that currently read the VTOC directly. Conversion of these programs to use CVAF (if that is necessary) may be a larger job than converting the VTOCs themselves. Plan to convert the programs to use the CVAF early in the migration period.

Performance Considerations

An indexed VTOC gives better performance through improved utilization of channels and devices. It efficiently allocates a new data set, renames an existing data set, and obtains a data set name. An indexed VTOC also provides significant improvements for allocating an existing data set, deleting a data set, using OPEN for input, and using OBTAIN by data set name.

Planning for User Programs That Read VTOCs

Programs that use DADSM macros to access nonindexed VTOCs continue to work on nonindexed and indexed VTOC volumes without any changes. However, programs that use other methods (for example, SAM or EXCP) to read the VTOC, or rely on the contents of certain VTOC fields, may require changes to run against volumes containing active indexed VTOCs. They may experience different performance characteristics when run on indexed VTOC volumes. The three VTOC fields that may cause problems are listed below, with an indication of which CVAF functions may be used to obtain the same information.

1. The Format-4 DSCB pointer to the highest Format-1 DSCB in the VTOC (DS4HPCHR): This field is used by some VTOC-reading programs to determine where to stop reading the VTOC, rather than allowing the reading program to continue to the end of the VTOC extent. When a VTOC is converted to indexed format, this field points to the end of the VTOC extent and causes the programs to read to the end of the extent anyway. Programs using this field work on indexed VTOC volumes—but take longer to run because they are reading through the preformatted Format-0 (empty) DSCBs at the end of the VTOC extent. The CVAF macro CVAFSEQ allows reading of indexed or nonindexed VTOCs in physical-sequential order (equivalent order to use of SAM), and allows reading of indexed VTOCs in alphabetic order by data set name. If physical order is specified, the read of Format-0 DSCBs at the end of the VTOC still occurs.
2. The Format-4 DSCB field contains the number of available DSCBs (DS4DSREC): This field may be used by a VTOC-reading program to indicate how full the VTOC is. When the VTOC is converted to indexed format, the field still contains a value but the value is no longer maintained and may not be accurate. Such programs should use the CVAFDSM macro with parameters COUNT = YES and MAP = VTOC. This form of the macro returns the number of available DSCBs in the VTOC.
3. The Format-5 DSCB free space fields (DS5AVEXT, DS5EXTAV, DS5MAVET): An indexed VTOC contains only one “dummy” Format-5 DSCB, and the available space fields in it contain zeros. Programs that use the LSPACE function of DADSM continue to work because LSPACE uses CVAF on an indexed VTOC supported system. Other programs that require free space information should use the CVAF macro CVAFDSM with the parameters COUNT = NO and MAP = VOLUME specified.

Determining Which VTOCs Should Be Indexed VTOCs

You must decide which VTOCs are to be in the indexed VTOC format. Two criteria should be considered:

- VTOC size (number of data sets in the VTOC)
- VTOC activity (amount and type)

Any VTOC larger than 2 tracks should be converted to the indexed format for faster access to Format-1 DSCBs via the index. Performance improves by avoiding the use of the chains of Format-5 DSCBs which a nonindexed VTOC may have (especially if free space is badly fragmented). Because DADSM tries to obtain a good match between the size of space available and space required, extensive Format-5 DSCB searches may result.

A frequently accessed VTOC has the most potential for performance improvement if it is indexed. However, the type of VTOC access is also an important factor.

An additional factor that may make a difference in some installations is the change in allocation technique that occurs with the use of indexed VTOC on a

volume. The old VTOC uses a "best fit" algorithm. (The routines select the smallest free extent that is big enough to satisfy the request for space on the volume, even if it is necessary to go through many Format-5 DSCBs to determine the point of best fit.) Indexed VTOC uses a "first fit" philosophy, placing an extent in the first location that is large enough to satisfy the request even if a closer fit exists further down the volume. Dependence on the old algorithm for certain volumes may be a factor in deciding which volumes to convert to indexed format.

The prime candidates for conversion to an indexed VTOC are the time sharing option (TSO) volumes, or any other volumes that have a high level of allocate/delete activity. Volumes containing data bases for online application systems can be converted later because they generally do not have heavy VTOC activity except at application initialization time. Volumes with VTOCs less than 3 tracks in size and volumes with seldom-used data may not be worth converting for performance reasons. The advantages of standardizing on one VTOC format in the installation may make conversion desirable.

When you decide which VTOCs are to be converted, make a list of volumes by priority. From this list, a migration schedule can be made. As with any product that is new to an installation, a gradual migration is advisable. Usually, the volumes are divided into logical groups according to user group within the priority. The migration plan for each group can be accomplished within a reasonable length of time so that one user group is not busy with the migration over a long period.

Before any VTOCs are actually converted, the results of the analysis of indexed VTOC impact on user and IBM programs must be considered and any outstanding problems resolved.

Establishing a VTOC Index at Volume Initialization Time

You may run the stand-alone or system version of Device Support Facilities, specifying the location and (optionally) the size of the VTOC and index on the INIT command. (*Device Support Facilities User's Guide and Reference* describes the syntax of the INIT command and includes some examples of its use.)

The VTOC index is usually placed next to the VTOC. The VTOC and index are good candidates for placement under a fixed head.

Although the VTOC location must be specified, it is possible to allow Device Support Facilities to calculate a size for the index based on the size of the VTOC. Then you do not need to get involved in complex calculations, which may be inadvisable for initial efforts with indexed VTOCs. If you prefer to do the calculations to estimate the correct size for a VTOC index, see "The VTOC Index" in *Device Support Facilities User's Guide and Reference*.

Building a VTOC Index on an Initialized Volume

An initialized volume already containing a VTOC is converted into an indexed VTOC format by using the system version of the BUILDIX command. The stand-alone version does not support the BUILDIX command. Before running Device Support Facilities, or its JCL, you must allocate the VTOC index data set. (You may use a DD statement in the same job step. For an example of a DD statement, see *Device Support Facilities User's Guide and Reference* in the section on BUILDIX). When using BUILDIX, it is necessary for the installation to specify the VTOC index size. "The VTOC Index" in *Device Support Facilities User's Guide and Reference* shows how to determine the correct size. The naming convention 'SYS1.VTOCIX.volser' is recommended. The name should be unique throughout the system in order to prevent ENQ lockouts.

The BUILDIX command is also used to convert a volume containing an indexed VTOC back to nonindexed format.

Compatibility

This section concentrates on the compatibility of DASD volumes when indexed VTOC volume support is installed.

Volume compatibility is a key feature of an indexed VTOC. It is possible to share an indexed or nonindexed VTOC volume between systems for concurrent access, or to move an indexed or nonindexed VTOC volume between systems. However, in each case there are a number of important considerations. The group of PTFs sometimes referred to as the "index bit reset feature" affects all cases. The group includes UX90053 (VS1) and UD19460 (DOS). These PTFs should be installed on any nonindexed VTOC-supported systems that may use an indexed VTOC volume. This code ensures that the VTOC index is disabled on nonsupported systems where it is appropriate. Failure to install these PTFs could result in the VTOC index becoming out of synchronization with the VTOC itself.

These PTFs provide support for invalidating VTOC indexes only where it is necessary. The code does not write records to SYS1.LOGREC to indicate that invalidation has occurred. When an index has been invalidated on a nonsupported system, subsequent use of the volume on a supported system might cause records to be written to SYS1.LOGREC, depending on specific circumstances. For example, if the invalidation causes the DOS bit to be zeroed but not the index bit (allocate or extend processing calling the DOS VTOC convert routine), subsequent processing on a supported system causes a LOGREC record to be written, an IEC606I message to appear, and a dump to be taken to a SYS1.DUMP data set. The only way to detect the condition is by running IEHLIST LISTVTOC against the suspect volume under a supported system. For further information on IEHLIST, see *Utilities*.

Note: The following discussion assumes that the PTFs are installed on all systems that do not contain support for indexed VTOC.

Sharing Volumes between Systems

Nonindexed VTOC volumes can still be shared between systems that include VTOC support (MVS, VS1, SVS, OS, and DOS), regardless of whether those systems contain indexed VTOC support.

For indexed VTOC volumes, support must be installed on all systems having access to that volume, particularly when sharing systems can all update the VTOC. But it is not advisable to attempt to share between supported and nonsupported systems even if the nonsupported system is not updating the VTOC. It is not possible for the installation to guarantee that a VTOC update cannot occur after the volume is accessible. The nonsupported system may invalidate the VTOC index and eliminate the performance advantages.

Invalidation of a VTOC Index

The conditions under which a VTOC index can be invalidated include:

- An error detected during processing causes invalidation by indexed VTOC-supported software
- Use of the VTOC on a nonsupported system or in circumstances that caused the DOS bit to be turned off in the Format-4 DSCB (for example, allocation of space on an indexed VTOC volume under a nonsupported system)
- Use of the Device Support Facilities BUILDIX command with the OSVTOC operand

The result in the first two cases is that the next request on a supported system causes conversion of the VTOC to OS format (Format-5 DSCBs will be rebuilt). In the third case, BUILDIX processing issues a dummy allocate so that the conversion occurs before BUILDIX processing is completed.

Moving Volumes between Systems

This category includes cases with volumes staying where they are located and the system changing. For example, an installation may have volumes that remain permanently resident on a particular configuration and different systems are brought up on that configuration at certain times. This may occur with fixed media DASD (for example, IBM 3350, 3375, and 3380).

For nonindexed VTOC volumes, nothing has changed. These volumes can be moved between systems regardless of whether those systems contain indexed VTOC support.

For indexed VTOC volumes where support of one of the PTFs referenced above is installed on all involved systems, DASD portability is available. However, certain types of volume movement and access may cause the VTOC index to be invalidated. When the volume is mounted on a supported system after invalidation under another system, the VTOC index on that volume is invalid. Processing continues without the use of the index. It is possible to detect this condition by running IEHLIST LISTVTOC with the INDEXDSN option

against the volume while it is mounted to a supported system. The IEHLIST output indicates a disabled index.

If the installation must move an indexed VTOC volume to a nonsupported system that does not contain one of the "index bit reset feature" PTFs, the Device Support Facilities BUILDIX function should be used to convert the VTOC to nonindexed format before moving the volume to the other system. This prevents the possibility of processing an incorrect VTOC index when the volume is later returned to the supported system, because this conversion process effectively invalidates the VTOC index. After the volume is returned, the BUILDIX function can be used again to rebuild the index based on the data that is in the VTOC at that time, including changes that occurred on the nonsupported system.

Estimating Resource Requirements

Space estimation guidelines for VTOC indexes are found in "The VTOC Index" in *Device Support Facilities User's Guide and Reference*. Virtual storage requirements in the private area for buffers for VTOC indexes range from about 10K to 12K bytes, depending on the function to be performed and the number of index levels in the VTOC.

Mass Storage System Considerations

Indexed VTOC software contains no special processing for MSS, and MSS volumes are treated as real volumes. An installation with a 3850 must decide whether to use indexed VTOC on virtual volumes, and, if so, where the VTOC index is to be placed. The only performance difference that alters the usual technique for deciding whether to use indexed VTOC on a virtual volume is the way initial access to the VTOC index is handled. If the index is on the first cylinder of the volume, there is no problem, because it is staged up along with the VTOC when the volume is mounted. However, some installations have adopted the practice of setting aside tracks 1 through 18 for the VTOC on virtual volumes. In these cases, a restructure of the volume is required to place the index on the first cylinder of the volume. If this is not done and the index is placed elsewhere on the first cartridge, the following sequence may occur when a request is made to allocate a new data set on the virtual volume:

1. Step initiation requests that MSS mount the volume.
2. MSS mounts the volume and stages up the first cylinder (containing the VTOC).
3. Volume verification invokes DADSM to allocate space on the volume for the new data set.
4. DADSM invokes CVAF, which recognizes that a control block structure for the VTOC index did not exist.

5. CVAF initializes the control block structure, and then performs the necessary I/O to the index.
6. This index I/O results in a cylinder fault.

At this point it is possible that the cartridge picked for the mount request is on the way back to the home cell. If the index is on the first cartridge, then the cylinder fault is not satisfied until the cartridge is returned to its home cell, repicked, and the required cylinder staged up. If the index is placed on the second cartridge, it may be possible to overlap the first cartridge's activity with that of the second cartridge, depending on the availability of more than one data recording device (DRD) to do the work. However, this results in allocation of two different pages for the VTOC and index with longer seek times.

If MSS volumes tend to stay mounted for extended periods, the extra activity caused by initialization of control blocks for the first request is not a significant overhead.

Security

Indexed VTOC security features parallel those of the VTOC to a large extent; however, the VTOC index has additional support. For example, it is possible to protect the VTOC index with a password but not the VTOC itself. The facilities available for providing security for VTOC indexes include:

- Password protection
- Resource Access Control Facility (RACF)
- Authorized program facility (APF)
- DADSM installation exits

Password Protection

The VTOC index can be protected as a data set with a password via the IEHPROGM ADD command or the PROTECT macro. Update-level protection stops OPEN of the VTOC index for output and SCRATCH or RENAME of the VTOC index. Read-level protection stops the VTOC index from being opened for input. This protection only affects OPEN, SCRATCH, and RENAME of the VTOC index as a data set; it does not affect DADSM/CVAF processing. If the VTOC index has password protection at the update level, a user job may still allocate new data sets and perform other VTOC index updates without providing the update password. Password protection prevents opening of the VTOC index as a data set. This may be worthwhile, because you can prevent VTOC index updates from occurring outside of DADSM/CVAF processing (for example, by service aid IMASPZAP).

Password protection for the VTOC itself is not available. If an attempt is made to change the VTOC using IMASPZAP, a message requesting permission to alter the VTOC is sent to the operator. This message does not appear if IMASPZAP is used to change the VTOC index, so password protection of the index is a method of getting similar checking for VTOC index updates from IMASPZAP.

Data set-level protection for user data sets in the VTOC and VTOC index provides some level of VTOC and VTOC index protection, because user data set passwords are required to SCRATCH, RENAME, or extend the data set. All these functions require updates to both the VTOC and its index.

Resource Access Control Facility (RACF)

It is possible to protect the VTOC itself with the DASDVOL/volser entity. RACF protection for the VTOC index is equivalent to this. It is not a data set-level protection; it is implemented via the DASDVOL/volser entity for both VTOC and VTOC index. With this level of protection in use, update-level authorization is required to open either the VTOC or the VTOC index for output. Alter-level authorization is required to SCRATCH or RENAME the VTOC index.

As with password protection, RACF protection applies only to OPEN, SCRATCH, and RENAME. CVAF and DADSM do not invoke RACF to check for passwords. Similarly, data set-level protection via RACF for user data sets in the VTOC and VTOC index provides some level of RACF protection for VTOC entries. For further information, see *RACF General Information Manual*.

Authorized Program Facility (APF)

The authorized program facility allows specification of authorized load modules and authorized program libraries. Regardless of the existence or nonexistence of password or RACF protection, any program that opens a VTOC or VTOC index for output or attempts to SCRATCH or RENAME a VTOC index must be an authorized load module residing in an authorized program library. IEHPROGM meets these requirements; disposition processing and TSO do not. If a VTOC index is active, it is not possible to SCRATCH or RENAME it even if the authorization requirements are met.

In a multiprocessor installation, if the environment is mixed indexed VTOC and nonindexed VTOC systems, the nonindexed VTOC systems can perform all the operations on a VTOC index that they can perform on a normal data set. For example, an unauthorized program can SCRATCH the index if it is not protected. This consideration, plus the fact that IEHPROGM can be used to SCRATCH the index on indexed VTOC or nonindexed VTOC systems, makes some form of protection for the index desirable.

DADSM Installation Exits

Password protection and RACF protection do not apply when VTOC updates are made via DADSM/CVAF. The DADSM installation exits may be used for rejecting attempts to perform operations that cause VTOC or VTOC index updates in which password and RACF protection is not invoked.

Backup and Recovery

Using indexed VTOCs introduces a few new considerations that affect backup and recovery. They are described below as they affect data set and volume backup and recovery.

Data Set Backup and Recovery

Use of an indexed VTOC does not affect the backup of user data sets. However, you should not back up the VTOC index itself as a data set. Restoring just the index without an equivalent restore of the VTOC may cause the two to be out of synchronization. The VTOC index can be rebuilt at any time, using the Device Support Facilities BUILDIX function, and there is no need to back up the VTOC index. Loss of the VTOC prevents rebuilding of the VTOC index.

Volume Backup and Recovery

When a volume containing a VTOC index is backed up with a total volume dump, the VTOC index is copied out along with the rest of the volume.

The only other consideration is that Data Facility Data Set Services (DFDSS), or a functional equivalent, should be used to do online restores of volumes containing indexed VTOCs if the location of the VTOC or the location or size of the index changes between backup and restore. Data Facility Data Set Services contains support for refreshing in storage control blocks containing data relating to VTOC and index location and size. For more information, see *DFDSS: User's Guide and Reference*.



Appendix A. Data Facility Product Features

The following tables summarize the features of MVS/XA Data Facility Product that are new or substantially different from those included in the MVS/XA DFP Version 1 and other previous program products. In addition, the tables list which commands and macros, if any, have been modified to support the new feature, and where information about that feature is documented.

Refer to the preface for the complete title and order number of each publication listed.

New Device Support

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
3262 Model 5 Printer	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Updating the UCS image table	IGGUCSIT macro	<i>System – Data Administration</i>
	Specifying UCS and FCB images	SETPRT macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>Checkpoint/Restart User's Guide</i>
	Creating an FCB module with IEBIMAGE	FCB statement	<i>Utilities</i> <i>System Messages</i>
3422 Magnetic Tape Subsystem	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
3430 Magnetic Tape Subsystem	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
3480 Magnetic Tape Subsystem	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
(3420 compatibility)	Specifying 3420 compatibility mode	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
(3480 full function)	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
	Using 18-track tapes	—	<i>Magnetic Tape Labels and File Structure Administration</i>
	Controlling message display	MSGDISP macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>System—Data Administration</i>
	Controlling buffering	SYNCDEV macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i>
	Using fast positioning during processing	NOTE macro POINT macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>System—Data Administration</i>
3880 Storage Control Model 21	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Error messages	—	<i>System Messages</i>
	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
	Controlling the cache	LISTDATA command SETCACHE command	<i>Cache Device Administration</i>
3880 Storage Control Model 23	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Error messages	—	<i>System Messages</i>
	Controlling the cache	LISTDATA command SETCACHE command	<i>Cache Device Administration</i>
4245 Printer	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
	Updating the UCS image table	IGGUCSIT macro	<i>System – Data Administration</i>
	Specifying UCS and FCB images	SETPRT macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>Checkpoint/Restart User's Guide</i>
4248 Printer	Description of programming support	—	<i>Data Facility Product Version 2: General Information</i>
	Specifying the device during system generation	This task now supported by MVS configuration program	<i>MVS Configuration Program</i>
	Updating the UCS image table	IGGUCSIT macro	<i>System – Data Administration</i>
	Specifying UCS and FCB images	SETPRT macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>Checkpoint/Restart User's Guide</i>
	Creating an FCB module with IEBIMAGE	FCB statement	<i>Utilities</i> <i>System Messages</i>

New Programming Support

MVS/XA Data Facility Product Version 2 Release 1.0

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
Catalog address space support	Establishing a system address space to be used by the catalog management function	—	<i>Catalog Administration Guide</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
DADSM ERASE support	Data security for scratch	—	<i>Data Administration Guide</i>
Erase on scratch security	RACF-controlled erasure	ERASE NOERASE parameter of ALTER, DEFINE ALTERNATEINDEX, DEFINE CLUSTER, DELETE, EXPORT, and IMPORT commands; DEFINE NONVSAM and DELETE NONVSAM commands	<i>Access Method Services Reference</i> <i>Catalog Administration Guide</i> <i>Data Administration Guide</i> <i>VSAM Administration Guide</i>
DADSM partial release function	Releasing unused space from a data set	PARTREL macro	<i>System – Data Administration</i>
DADSM REALLOC enhancements	Allocating absolute extents	REALLOC macro	<i>System – Data Administration</i>
Device console services	Testing all paths to a device and displaying information about the logical status of each path	DEVSERV command	<i>System Commands</i>
EXAMINE command	Collecting and analyzing information on the structural consistency of VSAM key-sequenced data set clusters	EXAMINE command	<i>Access Method Services Reference</i>
EXPORT and IMPORT by CIMODE	Exporting by control interval for entry-sequenced data set base cluster data records	CIMODE RECORDMODE parameter of EXPORT command; IMPORT command	<i>Access Method Services Reference</i> <i>System Messages</i> <i>VSAM Administration Guide</i>
IEBCOPY reblocking	Reblocking overlay modules; maximum output block size for load modules increased to 32760 bytes	ALTERMOD and COPYMOD functions of IEBCOPY	<i>Utilities</i> <i>Linkage Editor and Loader User's Guide</i>
ISO/ANSI/FIPS Version 3 label conversion	Converting non-Version 3 tape labels automatically to Version 3 tape labels, using WTOR installation exit	—	<i>Magnetic Tape Labels and File Structure Administration</i> <i>System Messages</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
RACF 1.7 tape data set protection	Protecting tape data sets, nonlabeled (NL) tapes, and tapes using bypass label processing (BLP)	—	<i>Data Administration Guide</i> <i>Magnetic Tape Labels and File Structure Administration</i> <i>System Messages</i>
Tape block count checking	Verifying the tape block count at close or end of volume	—	<i>Data Administration Guide</i>
VSAM JRNAD exits	Canceling an I/O request	—	<i>VSAM Administration Guide</i>
VSAM VERIFY support	Updating control blocks from the catalog	VERIFY macro ACTION=REFRESH parameter	<i>VSAM Administration: Macro Instruction Reference</i>
3480 label processing	Protecting tape data sets, nonlabeled (NL) tapes, and tapes using bypass label processing (BLP)	—	<i>Magnetic Tape Labels and File Structure Administration</i> <i>System—Data Administration</i>

MVS/XA Data Facility Product Version 2 Release 2.0

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
Interactive storage management facility (ISMF) data set application	Interactively analyzing, managing, and reporting on data and storage	—	<i>ISMF User's Guide</i> <i>Data Facility Product: Customization</i> <i>System—Data Administration</i> <i>System Generation</i>
CVAF filter (CVAFFILT)	Retrieving data set DSCBs from DASD VTOC into buffers provided by caller	CVAFFILT macro	<i>System—Data Administration</i>
Additional block sizes for integrated facility catalogs and for VSAM data sets cataloged in integrated catalog facility catalogs	Multiples of 512 bytes for blocks 512 through 8K bytes, and multiples of 2K bytes for blocks 8K through 32K bytes	CONTROL-INTERVALSIZE parameter of DEFINE ALTERNATEINDEX, DEFINE CLUSTER, and DEFINE USERCATALOG MASTERCATALOG	<i>Access Method Services Reference</i> <i>Catalog Administration Guide</i> <i>VSAM Administration Guide</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
Interactive storage management facility (ISMF) volume application	Interactively analyzing, managing, and reporting on volume storage	FILTER CLEAR and REFRESH commands	<i>ISMF User's Guide</i> <i>Data Facility Product: Customization</i> <i>System - Data Administration</i> <i>System Generation</i>
Enhanced integrated catalog facility catalog recovery	Recovering catalogs with minimal disruption to users	LOCK UNLOCK parameter for ALTER, DEFINE USERCATALOG MASTERCATALOG, and IMPORT commands	<i>Access Method Services Reference</i> <i>Catalog Administration Guide</i>
Communicating with catalog address space (CAS)	Communicating with CAS to directly control some recovery capabilities and obtain information for problem diagnosis	ABEND, CLOSE, ENTRY, LIST, OPEN, RESTART, VCLOSE options of MODIFY CATALOG command	<i>Catalog Administration Guide</i>
Integrated catalog facility EXPORT enhancements	Exporting aliases for an integrated catalog facility catalog; system management facilities (SMF) record written upon successful export of an integrated catalog facility catalog	ALIAS NOALIAS parameter for IMPORT command; system management facilities (SMF) WTM macro	<i>Access Method Services Reference</i> <i>Catalog Administration Guide</i>
Linear data sets	VSAM data sets with data in the order in which it was entered but with no control information	LINEAR keyword for DEFINE CLUSTER command; TYPE(LINEAR) keyword parameter for ALTER command; LDS option of TESTCB-ACB macro	<i>Access Method Services Reference</i> <i>VSAM Administration Guide</i> <i>VSAM Administration: Macro Instruction Reference</i>
Data set expiration beyond 1999	Maximum expiration date changed to last day of 2155; introduction of never scratch dates	TO and FOR parameters for all DEFINE commands except DEFINE ALIAS	<i>Access Method Services Reference</i> <i>Magnetic Tape Labels and File Structure Administration</i> <i>System - Data Administration</i> <i>VSAM Administration Guide</i>

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
Extended access of job file control block information	Allocating retrieval area for information from JFCB above 16 megabytes	RDJFCB function	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>System – Data Administration</i>
More than 1635 data definition statements	Coding up to 3273 DD statements for a job step	Task input/output table (TIOT)	—
Concatenation constraint relief	Concatenating partitioned data sets and some sequential data sets using less strict concatenation rules	—	—
Iogen restructure	Defining an I/O configuration	Iogen	<i>MVS Configuration Program</i>
VSAM alternate index synchronization	Base record reprocessed if an alternate index upgrade is interrupted	—	<i>VSAM Administration Guide</i>
VSAM improved read integrity	Improved handling of record-not-found situations for shared key-sequenced data sets	—	—
VSAM LSR pools for data components and index components	Improved handling of record-not-found situations for shared key-sequenced data sets	DATA INDEX added to TYPE parameter for BLDVRP macro; MACRF = LSR for ACB, GENCB-ACB, and MODCB-ACB macros	<i>VSAM Administration Guide</i> <i>VSAM Administration: Macro Instruction Reference</i>
Virtual storage constraint relief	VSAM, open/close/EOV, and DADSM modules moved from LPA to ELPA; VSAM control blocks and I/O buffers may reside above 16 megabytes virtual	MODE parameter for BLDVRP, CLOSE, DLVRP, and OPEN macros; RMODE31 parameter for ACB, BLDVRP, GENCB, and MODCB macros; LOC parameter for GENCB macro; SHRPOOL parameter for DLVRP macro; new DADSM LSPACE macro	<i>Data Administration Guide</i> <i>Data Administration: Macro Instruction Reference</i> <i>System – Data Administration</i> <i>VSAM Administration Guide</i> <i>VSAM Administration: Macro Instruction Reference</i>



Appendix B. Using VSAM in an MVS/XA Environment

VSAM has the following capabilities for users who run programs in 31-bit addressing mode. Installations planning to take advantage of 31-bit addressing capabilities for VSAM should be familiar with the information discussed in this appendix and in *VSAM Administration Guide*.

VSAM Record Management Code and Open/Close/EOV Code Reside in the ELPA

The majority of VSAM record management and open/close/EOV modules have been moved from the MVS link pack area (LPA) to the MVS/XA extended link pack area (ELPA). This change increases the amount of storage available to users below 16 megabytes virtual.

Applications residing below 16 megabytes virtual can access VSAM code in the ELPA with the help of an interface module. No changes to user programs are required.

VSAM I/O Buffers and Control Blocks Can Be Obtained above 16 Megabytes Virtual

User programs that run in 31-bit addressing mode can request and access VSAM I/O buffers and control blocks from any part of virtual storage, including the area above 16 megabytes virtual. User programs using VSAM move mode processing may run in 24-bit addressing mode with buffers or control blocks, or both, above 16 megabytes virtual but may not directly reference buffers, control blocks, or parameter lists included in control blocks.

To obtain the extended VSAM buffer and control block support, users must specify 31-bit buffers or 31-bit control blocks, or both, in the ACB at OPEN time.

Multiple LSR Pools Can Be Obtained in an Address Space

VSAM users running programs in 31-bit addressing mode can share local resources by building multiple local shared resource (LSR) pools in an address space. Up to 16 LSR data pools and 16 LSR index pools can be built for each address space, and I/O buffers and control blocks may reside above or below 16 megabytes virtual.

Restrictions

When using the extended 31-bit VSAM support, you should plan your activity so that you can accommodate the following restrictions:

- You cannot specify 31-bit VSAM support with JCL parameters. (For 31-bit VSAM support, you must code the RMODE31 parameter on the ACB and BLDVRP macros.)
- All VSAM control blocks that currently have fields defined as 31-bit addresses **must** contain complete 31-bit addresses. You may not use the high-order byte of a 31-bit address field as a user-defined flag field. This is true whether you are running in 24-bit or 31-bit addressing mode.
- A user exit routine that is not loaded by VSAM will be entered in the addressing mode of the program calling VSAM.
- To use 31-bit VSAM support, you must recompile that portion of your program that contains the ACB, BLDVRP, and DLVRP macro specifications and specify the appropriate 31-bit parameters.
- If the ACB or parameter list for OPEN/CLOSE and BLDVRP/DLVRP resides above 16 megabytes virtual, you must use the 31-bit form of these macros.

Detailed information on running in 31-bit addressing mode is contained in *VSAM Administration Guide*. For information on how to code VSAM macros in 31-bit addressing mode, see also *VSAM Administration: Macro Instruction Reference*.

Glossary of Terms and Abbreviations

The following terms are defined as they are used in this book. If you do not find the term you are looking for, see the index or the *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699.

ACB. Access method control block.

access method control block. A control block that links an application program to VSAM or ACF/VTAM.

access method services. A multi-function service program that defines VSAM data sets and catalogs and allocates space for them, converts indexed-sequential data sets to key-sequenced data sets with indexes, modifies data set attributes in the catalog, reorganizes data sets, facilitates data portability between operating systems, creates backup copies of data sets and indexes, helps make inaccessible data sets accessible, and lists the records of data sets and catalogs.

AIX. (*See* alternate index.)

alias. An alternative name for an entry.

alias entry. An entry that relates an alias (alternate entry name) to the real entry name of a user catalog or non-VSAM data set.

alternate index. An ordered collection of records, each consisting of a key (called the alternate key) and one or more pointers. An alternate index is used by VSAM to sequence and locate the records of a key-sequenced or entry-sequenced VSAM data set. An alternate index is organized as a key-sequenced data set. (*See also* alternate key, base cluster, and path.)

alternate index entry. A catalog entry that contains information about an alternate index. An alternate index is conceptually a key-sequenced cluster, and is cataloged in the same way. An alternate index entry points to a data entry and an index entry to describe the alternate index's components, and to a cluster entry to identify the alternate index's base cluster. (*See also* cluster entry.)

alternate index record. A collection of items used to sequence and locate one or more data records in a base cluster. Each alternate index record contains an

alternate key value and one or more pointers. When the alternate index supports a key-sequenced data set, each data record's prime key value is the pointer. When the alternate index supports an entry-sequenced data set, the data record's RBA value is the pointer. (*See also* alternate index, alternate key, base cluster, and key.)

alternate key. One or more characters within a data record, used to identify the data record or control its use. Unlike the prime key, the alternate key can identify more than one data record. (*See also* key and key field.)

APF. (*See* authorized program facility.)

application. As used in this publication, the use to which an access method is put or the end result that it serves; contrasted to the internal operation of the access method.

authorized program facility. A facility that permits the identification of programs that are authorized to use restricted functions.

backup data set. A copy that can be used to replace or reconstruct a damaged data set.

base cluster. The VSAM cluster whose data records are to be accessed through a path or AIX. Usually, a base cluster is the key-sequenced or entry-sequenced data set which an alternate index supports (that is, an alternate index is used by VSAM to sequence and locate the data records of a base cluster). (*See also* alternate index and path.)

basic catalog structure. The name of the actual catalog structure with the integrated catalog facility environment. (*See* integrated catalog facility.) Integrated catalog facility is composed of a BCS together with its related VVDSs (VSAM volume data sets).

BCS. (*See* basic catalog structure.)

CA. (*See* control area.)

catalog. (*See* master catalog and user catalog.)

catalog connector. A catalog entry, called either a user catalog entry or a catalog connector entry, in the master

catalog that points to a user catalog's volume (that is, it contains the volume serial number of the direct access volume that contains the user catalog).

CCHH. The cylinder head record that gives the DASD location.

cell. An occurrence of information such as passwords, volume information, or associations.

CI. (*See* control interval.)

cluster. A data component and an index component when data is key sequenced; a data component alone when data is entry sequenced.

cluster entry. A catalog entry that contains information about a key-sequenced or entry-sequenced VSAM cluster: ownership, cluster attributes, and the cluster's passwords and protection attributes. A key-sequenced cluster entry points to a data entry and an index entry. An entry-sequenced cluster entry points to a data entry.

component. The data portion or, for a key-sequenced cluster, alternate index, or VSAM catalog, the index portion or a VSAM object. In this book, the components of an object are usually referred to as the object's data component and index component. Also, the cluster, data, or index fields of a subrecord.

control area. A group of control intervals used as a unit for formatting a data set before adding records to it. Also, in a key-sequenced data set, the set of control intervals pointed to by a sequence-set index record; used by VSAM for distributing free space and for placing a sequence-set index record adjacent to its data.

control area split. The movement of the contents of some of the control intervals in a control area to a newly created control area to make possible the insertion or lengthening of a data record when a free control interval was needed and there was none in the original control area.

control interval. A fixed-length area of auxiliary-storage space in which VSAM stores records and distributes free space. It is the unit of information transmitted to or from auxiliary storage by VSAM.

control interval access. The retrieval and storage of a VSAM data set's contents, based on the RBA of a control interval (that is, the user's program processes a control interval, rather than a data record, as a logical entity).

control interval split. The movement of some of the stored records in a control interval to a free control interval to make possible the insertion or lengthening of a record that will not fit in the original control interval.

control volume. A volume that contains one or more indexes of the catalog.

CRA. Catalog recovery area.

CSA. Common service area.

CVAF. Common VTOC access facility

CVOL. (*See* control volume.)

DASD. (*See* direct access storage device.)

data component. That part of a VSAM data set, alternate index, or catalog that contains the object's data records.

data control block. A control block used by access method routines in storing and retrieving data.

data entry. A catalog entry that describes the data component of a cluster, alternate index, page spaces, or catalog. A data entry contains the data component's attributes, allocation and extent information, and statistics. A data entry for a cluster's or catalog's data component can also contain the data component's passwords and protection attributes.

data integrity. Preservation of data or programs for their intended purpose. As used in this publication, the safety of data from inadvertent destruction or alteration.

data record. A collection of items of information from the standpoint of its use in an application, as a user supplies it to VSAM for storage.

data security. Prevention of access to or use of data or programs without authorization. As used in this publication, the safety of data from unauthorized use, theft, or purposeful destruction.

data set. The major unit of data storage and retrieval in the operating system, consisting of data in a prescribed arrangement and described by control information to which the system has access. As used in this publication, a collection of fixed- or variable-length records in auxiliary storage, arranged by VSAM in key sequence or in entry sequence. (*See also* key-sequenced data set and entry-sequenced data set.)

data set control block. A data set label for a data set in direct access storage.

data space. A storage area defined in the volume table of contents of a direct access volume for the exclusive use of VSAM to store data sets, indexes, and catalogs.

DCB. (*See* data control block.)

direct access. The retrieval or storage of data by a reference to its location in a data set rather than relative to the previously retrieved or stored data.

direct access storage device. A device in which the access time is effectively independent of the location of the data.

DSCB. (*See* data set control block.)

dynamic allocation. The allocation of a data set or volume by the use of the data set name or volume serial number rather than by the use of information contained in a JCL statement.

entry. A collection of information about a cataloged object in a VSAM master or user catalog. Each entry resides in one or more 512-byte record.

entry name. A unique name for each component or object as it is identified in a catalog. The entry name is the same as the dsname in a DD statement that describes the object.

entry sequence. The order in which data records are physically arranged (according to ascending RBA) in auxiliary storage, without respect to their contents. (*Contrast* to key sequence.)

entry-sequenced data set. A data set whose records are loaded without respect to their contents, and whose RBAs cannot change. Records are retrieved and stored by addressed access, and new records are added at the end of the data set.

ESDS. (*See* entry-sequenced data set.)

exception. An abnormal condition such as an I/O error encountered in processing a data set.

extent. A continuous space allocated on a direct access storage volume reserved for a particular data space or data set. An extent of a data set contains a whole number of control areas.

field. In a record or a control block, a specified area used for a particular category of data or control information.

GDG. (*See* generation data group entry.)

generation data group entry. An entry that permits non-VSAM data sets to be associated with other non-VSAM data sets as generation data sets.

generation data set. One of a collection of historically related non-VSAM data sets; the collection of these data sets is known as a generation data group.

generic key. A high-order portion of a key, containing characters that identify those records that are significant

for a certain application. For example, it might be desirable to retrieve all records whose keys begin with the generic key AB, regardless of the full key values.

index. As used in this publication, an ordered collection of pairs, each consisting of a key and a pointer, used by VSAM to sequence and locate the records of a key-sequenced data set; organized in levels of index records. (*See also* index level, index set, and sequence set.)

index component. That part of a key-sequenced data set, catalog, or alternate index that establishes the sequence of the data records within the object it indexes. The index is used to locate each record in the object's data component based on the record's key value.

index entry. A catalog entry that describes the index component of a key-sequenced cluster, alternate index, or catalog. An index entry contains the index component's attributes, passwords and protection attributes, allocation and extent information, and statistics.

index level. A set of index records that order and give the location of records in the next lower level, or of control intervals in the data set that it controls.

index record. A collection of index entries that are retrieved and stored as a group. (*Contrast* to data record.)

index set. The set of index levels above the sequence set. The index set and the sequence set together comprise the index.

integrated catalog facility. The name of the catalog that is a functional replacement for OS CVOLs and VSAM catalogs.

integrity. (*See* data integrity.)

internal sort. Sorting of data records into a new sequence using virtual storage and no temporary data sets on a direct access storage device.

IPL. Initial program load.

ISAM interface. A set of routines that allow a processing program coded to use ISAM (indexed sequential access method) to gain access to a key-sequenced data set with an index.

ISMF. Interactive storage management facility.

K. When referring to storage capacity, 2 to the 10th power; 1024 in decimal notation.

key. One or more characters within an item of data that are used to identify it or control its use. As used in this publication, one or more consecutive characters

taken from a data record, used to identify the record and establish its order with respect to other records. (See also key field and generic key.)

key field. A field located in the same position in each record of a data set whose contents are used for the key of a record.

key range. In VSAM, a particular key range (for example, A through F) that is specifically associated with one or more control ranges of a data set.

key sequence. The collating sequence of data records, determined by the value of the key field in each of the data records. May be the same as, or different from, the entry sequence of the records.

key-sequenced data set. A data set whose records are loaded in key sequence and controlled by an index. Records are retrieved and stored by keyed access or by addressed access, and new records are inserted in the data set in key sequence by means of distributed free space. RBAs of records can change.

KSDS. (See key-sequenced data set.)

linear data set. A VSAM data set that contains data but does not contain control information. A linear data set can be accessed as a byte-addressable string in virtual storage.

mass storage volume. The unit of mass storage in the 3850 Mass Storage System.

master catalog. A key-sequenced data set with an index containing extensive data set and volume information that VSAM requires to locate data sets, to allocate and deallocate storage space, to verify the authorization of a program or operator to gain access to a data set, and to accumulate usage statistics for data sets.

non-VSAM entry. A catalog entry that describes a non-VSAM data set. A non-VSAM entry contains the data set's volume serial number and device type. If the data set resides on a magnetic tape volume, the entry can also identify the data set's file number. When the data set resides on a direct access device, the operating system obtains further information by examining the data set's DSCB (Data Set Control Block) in the volume's VTOC (volume table of contents).

object. A logical entity created by VSAM, such as a cluster (VSAM data set) and its components, an alternate index and its components, a VSAM catalog and its components, a path, or a VSAM data space.

OS CVOL. (See control volume.)

password. A unique string of characters stored in a catalog that a program or a computer operator at the

console must supply to meet security requirements before the program gains access to a data set.

path. A data set name for the combination of an alternate index and its base cluster, or an alias for a VSAM data set.

path entry. A catalog entry that contains information about a path, and that points to the path's related objects.

physical record. On a track of a direct access storage device, the space between interrecord gaps.

pointer. An address or other indication of location. For example, an RBA is a pointer that gives the relative location of a data record or a control interval in the data set to which it belongs.

portability. The ability to use VSAM data sets with different operating systems. Volumes whose data sets are cataloged in a user catalog can be demounted from storage devices of one system, moved to another system, and mounted on storage devices of that system. Individual data sets can be transported between operating systems using access method services.

primary space allocation. Initially allocated space on a direct access storage device, occupied by or reserved for a particular data set. (See also secondary space allocation.)

program temporary fix. A temporary solution or by pass of a program diagnosed by IBM field engineering as the result of a defect in a current unaltered release of the program.

PTF. (See program temporary fix.)

RACF. Resource Access Control Facility.

RBA. (See relative byte address.)

record. (See index record or data record.)

recoverable catalog. A catalog defined with the recoverable attribute. Duplicate catalog entries are put into CRAs that can be used to recover data in the event of catalog failure. (See also CRA.)

relative byte address. The displacement of a data record or a control interval from the beginning of the data set to which it belongs; independent of the manner in which the data set is stored.

relative record. A data record whose position depends on its placement within a group of data records; its position, or record number, is its displacement, in records, from the beginning of the data set.

relative record data set. A data set whose records are loaded into fixed-length slots.

relative record number. A number that identifies not only the slot or record space in a relative record data set, but also the record occupying the slot.

reusable data set. A VSAM data set that can be used as a workfile regardless of its old contents.

RMF. Resource management facility.

RPL. Request parameter list.

RPL string. A set of chained RPLs (the set may contain one or more RPLs) used to gain access to a VSAM data set by action macros (GET, PUT, etc.). Two or more RPL strings may be used for concurrent direct or sequential requests made from a processing program or its subtasks.

RRDS. (*See* relative record data set.)

secondary space allocation. A contiguous space on a direct access device, occupied by or reserved for a particular data set, which is allocated after space in the primary extent has been exhausted. (*See also* primary space allocation.)

security. (*See* data security.)

sequence set. The lowest level of the index of a key-sequenced data set; it gives the locations of the control intervals in the data set and orders them by the key sequence of the data records they contain. The sequence set and the index set together comprise the index.

sequential access. The retrieval or storage of a data record in either its entry sequence or its key sequence, relative to the previously retrieved or stored record.

shared resource. A set of functions that points to the sharing of a pool of I/O related control blocks, channel programs, and buffers among several VSAM data sets open at the same time.

skip sequential access. Keyed sequential retrieval or storage of records here and there throughout a data set, skipping automatically to the desired record or collating position for insertion: VSAM scans the sequence set to find a record or a collating position.

slot. The space for a data record in a relative record data set.

SMF. System Management Facility.

source catalog. An existing catalog that may be exported into a target catalog.

spanned record. A logical record whose length exceeds control interval length, and crosses (or spans) one or more control interval boundaries within a control area.

sphere record. A collection of logically related subrecords in one VSAM logical record.

stage. (verb) To transmit data from a mass storage volume to a direct access storage staging drive.

subrecord. The user definition level of a sphere, such as an AIX, cluster, or generation data set.

target catalog. The catalog that data sets or a source catalog are imported into.

TSO. Time sharing option.

TTR. Track record.

user catalog. A catalog used in the same way as the master catalog, but optional and pointed to by the master catalog, and also used to lessen the contention for the master catalog and to facilitate volume portability.

user catalog connector. (*See* catalog connector.)

VSAM volume control record. The first logical record in the VVDS that contains information to manage DASD space and the BCS back pointers.

VSAM volume data set. The VSAM volume data set is used to describe data set characteristics of VSAM data sets residing on a given volume. There is one, and only one, VVDS for each volume containing VSAM data sets cataloged in an integrated catalog facility catalog.

VSAM volume record. The VSAM volume record is a VSAM logical record within a VVDS.

VTOC. Volume table of contents.

VVCR. (*See* VSAM volume control record.)

VVDS. (*See* VSAM volume data set.)

VVR. (*See* VSAM volume record.)



Index

A

Access Method Services Cryptographic Option
replaced by Data Facility Product 15

ACF/TCAM (Advanced Communications Function for
TCAM)
required level for Data Facility Product 11

ACF/VTAM (Advanced Communications Function for
VTAM)
required level for Data Facility Product 11

alias 42

allocation, secondary
controlling control area size of the integrated catalog
facility catalog 52
setting up the BCS 68
setting up the VVDS 66

ALTER command
constructing the integrated catalog facility
catalog 51
controlling integrated catalog facility catalog
performance options 47
structuring the integrated catalog facility catalog 54
with LOCK and UNLOCK 81

alter-level authorization 57, 100

alternate index 41

alternate master catalog
catalog conventions 46
configuration 48
converting VSAM master catalog 73
generating 42
in recovering a master catalog 83

AMASPZAP VERIFY 71

APF (authorized program facility) 100

Assembler H Version 2
required level for Data Facility Product 9

authorizing ISMF functions 26

B

backing up
catalogs 66, 84
data spaces 66
OS CVOLs 66
VSAM data sets 66

basic catalog structure
See BCS

BCS (basic catalog structure)
See also integrated catalog facility catalog
contents 39, 67-72
key-sequenced data set 40
recovering 77
reorganizing 56, 70, 71
setting up 67

size 67-70

BTAM/SP (BTAM/System Product)
required level for Data Facility Product 11

BUILDIX command 93-97, 101

C

CAMLST parameter 73

catalog
conversion 70
name retention 70

catalog, integrated catalog facility
backup 77, 84
configuration 48
constructing 51
conventions 46
entries 41
locking 81
procedures 54
recovery 77

catalog, VSAM
nonrecoverable 65
recoverable 65

CICS/VS (Customer Information Control System/Virtual
Storage)
required level for Data Facility Product 11

cluster 41

CNVTCAT command
converting a VSAM catalog into an integrated
catalog facility catalog 64-72

common VTOC access facility
See CVAF

compatible licensed programs
with Data Facility Product 11

control area size 68

control area split 69, 70

control interval size 68

control interval split 69, 70

conversion
a VSAM catalog to an integrated catalog facility
catalog 70
an application, user or data set 64
an OS CVOL to an integrated catalog facility
catalog 70
catalog 70
one catalog at a time 63
one volume at a time 63
using CNVTCAT 70

corequisite licensed programs
for Data Facility Product 18

Cryptographic Unit Support
required level for Data Facility Product 12

customization restructure 32

CVAF (common VTOC access facility) 91

CVAFDSM macro 94

CVAFSEQ macro 94

D

DADSM (direct access device storage management)
accessing nonindexed VTOC 93
accessing VTOC 89
installation exits 100
LSPACE function 94
Mass Storage System considerations 98
matching space available with space required 94
password protection of VTOC index 99-100
space management for integrated catalog facility 41, 53

DASD (direct access storage device)
compatible with Data Facility Product 5
data error detection and diagnosis 77
Data Facility Product
See DFP

DEFINE command
cataloging entries 41
controlling integrated catalog facility catalog performance options 47
in updating recovered catalog 80
rebuilding alias entries after conversion to integrated catalog facility 70, 75
recovering the BCS 53, 55
reducing maximum record size for the BCS 69
reorganizing the BCS 56
space required when building entries in the BCS 70
with LOCK and UNLOCK 81

DELETE command
in updating recovered catalog 80
recovering the BCS 55

determining which VTOCs should be indexed 94

devices
compatible with Data Facility Product 4
in compatibility mode 6
input/output 5

DFDS (Data Facility Device Support)
replaced by Data Facility Product 4, 15

DFDSS (Data Facility Data Set Services)
dumping non-VSAM data sets 66
in catalog backup 84
in data error diagnosis 78
required for installing DFP 10
substitute for DRWDASDR 16
substitute for IBCDMPRS 16
substitute for IBHDASDR 16

DFEF (Data Facility Extended Function)
replaced by Data Facility Product 4, 15

DFHSM (Data Facility Hierarchical Storage Manager)
in catalog backup 84
in catalog recovery 78
required level for Data Facility Product 11

DFP (Data Facility Product)
corequisite licensed programs 18
incompatible licensed programs 15

licensed programs compatible with 11
licensed programs required for 9

DFSORT (Data Facility Sort)
required level for Data Facility Product 11

DIAGNOSE command
locating discrepancies between the BCS and VSAM data sets 55, 80
verifying conversion to integrated catalog facility 71
VSAM data set recovery time 49

diagnosis restructure 32

distribution data sets 22

distribution libraries
See DLIBs

DLIBs (distribution libraries) 22
load module library 22
message library 22
panel library 22
placement 23
cost 23
performance 23
recovery 23
space allocation 23
table library 22

DOS bit
off 96, 97
on 90

DSF (Device Support Facilities) 9, 16

E

enhanced integrated catalog facility catalog recovery
catalog locking 81
considerations in a shared environment 82
description 80
restrictions 82
security 81

entry
in a catalog 41

erasing residual data 57

EREP (Environmental Recording, Editing, and Printing) 9

EXAMINE command 77

execution data sets 23
placement 24
cost 24
performance 24
recovery 24
space allocation 23

EXPORT command
and export of aliases 79
backing up data and associated catalog entries 66
backing up the integrated catalog facility catalog 54
moving VSAM data sets when volume is full 65
removing user catalog connector 70
reorganizing the BCS 56, 70
verifying conversion to integrated catalog facility 71

EXPORTRA command

not supported by integrated catalog facility 72
releasing space for conversion to integrated catalog
facility 65

F

free space
fragmented 94
information in VTOC index 89, 94
preparing for CNVTCAT processing 65
reorganizing the BCS after conversion to integrated
catalog facility 71
full volume 65
full VTOC 65

G

generation data group 42
glossary 113

I

I/O resources 59
IBM licensed programs
compatible with Data Facility Product 11
required before installing Data Facility Product 9
IEHLIST program
detecting invalidation of VTOC index 97
used to manage OS CVOL 75
verifying conversion to integrated catalog facility 71
IEHMOVE program
used to manage OS CVOL 75
IEHPROGM program
meets requirements of authorized program
facility 100
protecting the VTOC index 99
restricted to selected users 54
used to manage OS CVOL 75
IGG.CATLOCK facility 81
IMASPZAP program 99
IMPORT command
backing up the integrated catalog facility
catalog 55
moving VSAM data sets when volume is full 65
providing free space 71
reorganizing the BCS 56, 70
verifying conversion to integrated catalog facility 71
with ALIAS 79
with LOCK and UNLOCK 79, 81
IMPORTRA command
not supported by integrated catalog facility 72
releasing space for conversion to integrated catalog
facility 65

IMS/VS (Information Management System/Virtual
Storage)
required level for Data Facility Product 11
incompatible licensed programs
with Data Facility Product 15
index bit 90, 96
index bit reset feature 98
INIT command 95
initializing DASD volumes 19
installations with VSAM catalogs 73
installing Data Facility Product 3, 17
integrated catalog facility catalog
See also BCS, VVDS
catalog
backup and recovery 72, 77
configuration 48
constructing 51
conventions 46
locking 81
procedures 54
configuration 48
constructing 51
conventions 46
conversion
choosing a technique 63
CNVTCAT prerequisites 64-66
scheduling 60
testing 73
verifying 71
data set 41
backup and recovery 72
procedures 52
entries 41
I/O resources 59
installation tasks 60
procedures 54
space management 41, 59
system resources 58
virtual storage requirement 58
interactive storage management facility
See ISMF
ISMF (interactive storage management facility)
and DFDSS 10
and TSO users 24
authorization 26
data sets used 22
distribution data sets 22
DLIBs 22
placement 23
space allocation 23
execution data sets 23
placement 24
space allocation 23
general description 21
placement for load modules 25
required products 21
TLIBs 22
ISPF (Interactive System Productivity Facility) 10
table libraries 24
ISPTABL 24
ISPTLIB 25

J

- JES2 (Job Entry Subsystem)
 - with Data Facility Product 7
- JES3 (Job Entry Subsystem)
 - with Data Facility Product 7
- JOBCAT DD statement
 - conventions for data sets and integrated catalog facility catalogs 46, 48
 - revising data set procedures for integrated catalog facility catalogs 53, 72

L

- LISTCAT command
 - for data error diagnosis 77
 - rebuilding alias entries after conversion to integrated catalog facility 70, 75
- LISTCRA command
 - not supported by integrated catalog facility 72
- LISTVTOC command 71
- load module
 - ISMF 25
- load module library 22
 - as DLIB 22
 - as execution data set 23
- locking, catalog 81

M

- macros, VSAM
 - conversion to integrated catalog facility 73
- Mass Storage System
 - See MSS
- master catalog
 - alternate 73, 83
 - converting VSAM to integrated catalog facility 73
- message library 22
 - as DLIB 22
 - as execution data set 23
- MSS (Mass Storage System)
 - integrated catalog facility catalog considerations 48, 54
 - use with indexed VTOC 98
- MVS configuration program 18
- MVS/SP (MVS/System Product)
 - with Data Facility Product 7
- MVS/370
 - converting to MVS/XA 4

N

- non-VSAM data set 41

O

- Offline IBM 3800 Utility
 - replaced by Data Facility Product 15
- OPEN macro
 - protecting the VTOC index 99, 100
 - releasing multiple extents 54
- operating systems
 - compatible with Data Facility Product 3
- OS CVOL
 - converting to integrated catalog facility catalog 74
- OS/VS2 MVS
 - as a base for Data Facility Product 4

P

- page space 41
- panel library 22
 - as DLIB 22
 - as execution data set 23
- parameters not supported by integrated catalog facility 72
- password
 - VSAM, supported under integrated catalog facility 56
- path 41
- planning
 - for installation of Data Facility Product 3
- planning for conversion
 - of a VSAM catalog 63
 - of a VSAM master catalog 73
 - of an OS CVOL 74
 - to an integrated catalog facility catalog 63
- prerequisite licensed programs
 - for Data Facility Product 9
- PRINT command 73
- Print Services Facility
 - See PSF
- procedures
 - for catalogs 54
 - for data sets 52
- processing units
 - compatible with Data Facility Product 5
- Programmed Cryptographic Facility
 - required level for Data Facility Product 12
- PROTECT macro 99
- protecting ISMF functions 26
- PSF (Print Services Facility)
 - required level for Data Facility Product 13

PTF (program temporary fix) 96
publications
for Data Facility Product 27

R

RACF (Resource Access Control Facility)
alter-level authorization 57, 100
generic profiles 56
protection
for data set names 46
for integrated catalog facility catalog 46
integrated catalog facility catalog and associated entries 55, 56
VSAM and non-VSAM data sets 56
VTOC index 100
required level for Data Facility Product 11, 12
recovery
of catalog entries 78, 83
of master catalog 83
procedure 79
RENAME macro 99, 100
REPRO command
backing up data and associated VSAM catalog entries 66
converting an OS CVOL into an integrated catalog facility catalog 74
merging integrated catalog facility catalogs 56
not used for backing up the VVDS 56
releasing space for conversion to integrated catalog facility 66
reorganizing the integrated catalog facility catalog structure 64
RESERVE macro 91
Resource Access Control Facility
See RACF
RESTORE command 79
restructure
customization 32
diagnosis 32
revising procedures for catalogs 72
RMF (Resource Measurement Facility)
required level for Data Facility Product 11

S

SAM-E (Sequential Access Method-Extended)
replaced by Data Facility Product 15
SCRATCH macro 99, 100
secondary allocation
controlling control area size of the integrated catalog facility catalog 52

setting up the BCS 68
setting up the VVDS 66
security 56
setting up the BCS 67
setting up the VVDS 66
SMF (system management facilities)
processing data set extensions 53
recovering the integrated catalog facility catalog 49, 55
updating a recovered catalog 79
SMP (System Modification Program) 9, 18
SMP/E (System Modification Program Extended) 9, 18
space map 91
STPCAT DD statement
conventions for data sets and integrated catalog facility catalogs 46, 48
revising data set procedures for integrated catalog facility catalogs 53, 72
Storage Management Library 33
storage requirements
for Data Facility Data Set Services (DFDSS) 20
for Data Facility Product 17
for Device Support Facilities 19
for integrated catalog facility catalogs 58
for JES2 or JES3 18
for SMP or SMP/E 19
Storage Subsystem Library 34
SVC26 73
SYSCATLG member 73
SYSCATnn member 42, 73
system generation
requirements 17
system management facilities
See SMF
system resources 58

T

table library 22
as DLIB 22
as execution data set 23
ISPF 24
target libraries
See execution data sets
temporary data sets
See TLIBs
TLIBs (temporary data sets) 22
toleration PTFs 10, 13
TSO (time sharing option)
conversion to indexed VTOC 95
conversion to integrated catalog facility 60, 70
making ISMF available 24
TSO/E (Time Sharing Option Extensions) 11

U

UNIQUE parameter
for data sets 65
user catalog 41

V

VATLIST (volume attribute list) 48
VERIFY command
verifying conversion to integrated catalog facility 71
verifying the conversion 71
VIER (VTOC index entry record) 91
virtual storage access method
See VSAM
virtual storage requirements
for Data Facility Product 17
for IBM licensed programs 17
for integrated catalog facility catalogs 58
virtual volumes 48, 98
VIXM (VTOC index map) 91
VMDS (VTOC map of DSCBs) 91
volume compatibility 96
volume table of contents
See VTOC
volume, full 65
volumes
initializing 19
VPSM (VTOC pack space map) 91
VSAM (virtual storage access method)
catalog
backup and recovery 66
nonrecoverable 65
recoverable 65
data set
backup and recovery 66
macros
conversion to integrated catalog facility 73
VSAM catalog
nonrecoverable 65
recoverable 65
VSAM volume control record
See VVCR
VSAM volume data set
See VVDS
VSAM volume record

See VVR

VSCR (virtual storage constraint relief) 58
VTOC (volume table of contents)
backup and recovery
of data set 101
of volume 101
compatibility
moving volumes between systems 97
sharing volumes between systems 97
CVAF (common VTOC access facility) 91
full 66
illustration 90
Mass Storage System considerations 98
resource requirements 98
security
APF (authorized program facility) 100
DADSM installation exits 100
password protection 99
RACF 100
user programs which read VTOCs 93
VTOC index
building on an initialized volume 96
contents 90
determining which VTOCs should be indexed 94
establishing at volume initialization time 95
invalidation of 97
relationship between VTOC and its index 90
space map 91
structure 89-91
VTOC index map (VIXM) 91
VTOC map of DSCBs (VMDS) 91
VTOC pack space map (VPSM) 91
VVCR (VSAM volume control record) 40, 67
VVDS (VSAM volume data set)
See also integrated catalog facility catalog
contents 39, 66
entry-sequenced data set 40
setting up 66
VVR (VSAM volume record) 41

Numerics

31-bit addressing mode
VSAM support 111
3480 Magnetic Tape Subsystem
in compatibility mode 6
3800 Printing Subsystem Model 3
in compatibility mode 6

MVS/XA Data Facility Product Version 2:
Planning Guide
GC26-4147-2

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 automated 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 stamp 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 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 PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
P.O. Box 50020
Programming Publishing
San Jose, California 95150



Fold and tape

Please do not staple

Fold and tape



MVS/XA Data Facility Product Version 2:
Planning Guide
GC26-4147-2

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 automated 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 stamp 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 directly to the address in the Edition Notice on the back of the title page.)

GC26-4147-2

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 PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
P.O. Box 50020
Programming Publishing
San Jose, California 95150



Fold and tape

Please do not staple

Fold and tape





MVS/Extended Architecture
Data Facility Product Version 2:
Planning Guide

File Number S370-34

GC26-4147-02



Printed in U.S.A.