IBM

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

DFP
AMODE
DFP 31-bit
MVS/XA RM(
24-bit

IBM

MVS/Extended Architecture
Data Facility Product:
Planning Guide

Licensed
Program

Preface

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

- Installation of Data Facility Product
- Installation of or conversion to integrated catalog facility catalogs
- Conversion to indexed volume tables of contents (VTOCs)

Organization

This publication contains the following major parts:

- "Part 1: Planning for Data Facility Product Installation" on page 1, 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 licensed programs 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 15, provides guidelines for estimating storage for Data Facility Product and corequisite licensed programs, and indicates how to plan for the system generation process.

  - Chapter 3, "Updating the Publications Library" on page 19, 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" on page 25, explains the integrated catalog facility catalog environment and describes how to plan for catalog installation or conversion.

  - Chapter 4, "Integrated Catalog Facility Catalog Overview" on page 27, provides a description of the integrated catalog facility catalog and its environment.

  - Chapter 5, "Integrated Catalog Facility Catalog Installation" on page 31, describes the installation process and the decisions that must be made before installation.

- Chapter 6, "Integrated Catalog Facility Catalog Conversion" on page 49, describes the conversion of OS CVOLs or VSAM catalogs to integrated catalog facility catalogs and indicates how you should prepare for catalog conversion.

• "Part 3. Planning for the Indexed VTOC" on page 63, describes the structure of the indexed VTOC and outlines how to plan for VTOC conversion.

- Chapter 7, "Indexed VTOC Overview" on page 65, explains how the indexed VTOC is used.

- Chapter 8, "Indexed VTOC Conversion" on page 69, 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 79, lists the features of Data Facility Product that are new or substantially different from previous IBM licensed programs, and indicates where information on these subjects is documented.

• Appendix B, "Using VSAM in an MVS/XA Environment" on page 83, 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 85, defines the terms and abbreviations used 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 Catalog Administration Guide*, GC26-4041, describes data set catalogs in greater detail.

• *MVS/Extended Architecture Installation: System Generation*, GC26-4009, describes the process of system generation and how it relates to installation.

- *MVS/Extended Architecture System – Data Administration*, GC26-4010, describes the VTOC in greater detail.

- *MVS/Extended Architecture VSAM Administration Guide*, GC26-4015, describes VSAM and how to use VSAM data sets.

Related Publications

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

Short Title	Publication Title	Order Number
Access Method Services Reference	*MVS/Extended Architecture Integrated Catalog Administration: Access Method Services Reference*	GC26-4019
	MVS/Extended Architecture VSAM Catalog Administration: Access Method Services Reference	GC26-4075
Assembler H Version 2 General Information	*Assembler H Version 2 General Information*	GC26-4035
Cache Device Administration	*MVS/Extended Architecture Cache Device Administration*	GC26-4017
Catalog Administration Guide	*MVS/Extended Architecture Catalog Administration Guide*	GC26-4041
Checkpoint/ Restart	*MVS/Extended Architecture Checkpoint/Restart User's Guide*	GC26-4012
Conversion Notebook	*MVS/Extended Architecture Conversion Notebook*	GC28-1143
Cryptographic Unit Support General Information	*OS/VS1 and OS/VS2 MVS Cryptographic Unit Support General Information*	GC28-1015
CVAF Diagnosis Reference	*MVS/Extended Architecture Common VTOC Access Facility Diagnosis Reference*	SY26-3929
DADSM and CVAF Diagnosis Guide	*MVS/Extended Architecture DADSM and Common VTOC Access Facility Diagnosis Guide*	SY26-3896
DADSM Diagnosis Reference	*MVS/Extended Architecture DADSM Diagnosis Reference*	SY26-3904
DASD Migration Aid General Information	*DASD Migration Aid General Information*	GC26-3972
Data Administration Guide	*MVS/Extended Architecture Data Administration Guide*	GC26-4013

Short Title	Publication Title	Order Number
Data Administration: Macro Instruction Reference	*MVS/Extended Architecture Data Administration: Macro Instruction Reference*	GC26-4014
Data Facility Product: General Information	*MVS/Extended Architecture Data Facility Product: General Information*	GC26-4007
Device Support Facilities User's Guide and Reference	*Device Support Facilities User's Guide and Reference*	GC35-0033
DFDSS: General Information	*Data Facility Data Set Services: General Information*	GC26-3947
DFDSS: User's Guide and Reference	*Data Facility Data Set Services: User's Guide and Reference*	GC26-3949
EREP User's Guide and Reference	*Environmental Recording, Editing, and Printing (EREP) User's Guide and Reference*	GC28-1178
Global Resource Serialization	*OS/VS2 MVS Planning: Global Resource Serialization*	GC28-1062
IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility	*IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility*	SH35-0051
IBM 3800 Printing Subsystem Programmer's Guide	*IBM 3800 Printing Subsystem Programmer's Guide*	GC26-3846
IOCP User's Guide and Reference	*Input/Output Configuration Program User's Guide and Reference*	GC28-1027
JCL	*MVS/Extended Architecture JCL*	GC28-1148
Magnetic Tape Labels and File Structure Administration	*MVS/Extended Architecture Magnetic Tape Labels and File Structure Administration*	GC26-4003
MVS/SP General Information Manual	*MVS/System Product Version 2 General Information Manual*	GC28-1118
RACF General Information Manual	*Resource Access Control Facility (RACF): General Information Manual*	GC28-0722
SMP System Programmer's Guide	*System Modification Program (SMP) System Programmer's Guide*	GC28-0673
SMP/E General Information Manual	*System Modification Program Extended (SMP/E) General Information Manual*	GC28-1106
SMP/E User's Guide	*System Modification Program Extended (SMP/E) User's Guide*	SC28-1302

Short Title	Publication Title	Order Number
System – Data Administration	*MVS/Extended Architecture System – Data Administration*	GC26-4010
System Generation	*MVS/Extended Architecture Installation: System Generation*	GC26-4009
System Messages	*MVS/Extended Architecture Message Library: System Messages*, Volumes 1 and 2	GC28-1376 and GC28-1377
Utilities	*MVS/Extended Architecture Data Administration: Utilities*	GC26-4018
VSAM Administration Guide	*MVS/Extended Architecture VSAM Administration Guide*	GC26-4015
VSAM Administration: Macro Instruction Reference	*MVS/Extended Architecture VSAM Administration: Macro Instruction Reference*	GC26-4016

Summary of Changes

Release 1.2 Library Update, January 1987

New Device Support

Support has been added for the 3880 Model 23.

Support has been added for the 3480 Magnetic Tape Subsystem.

Support has been added for the 3880 Model 21.

Service Changes

Information has been added, corrected, or deleted to reflect technical service changes.

Release 1.2, February 1984

Documentation Additions and Changes

The title of this manual has been changed from *MVS/Extended Architecture Data Facilities Planning Guide* to *MVS/Extended Architecture Data Facility Product: Planning Guide*, in order to more accurately reflect the contents of the book. The order number (GC26-4040) remains the same.

Part 1, "Planning for Data Facility Product Installation, " has been added to this manual. Chapters 1 through 3 contain procedures for product and installation planning, including information on hardware and software requirements, storage requirements, and new publications. The succeeding parts and chapters have been renumbered.

The section on "Virtual Storage Access Method," formerly Part 3 of this manual, has been deleted. Information about planning for the 31-bit addressing capabilities of VSAM is contained in Appendix B, "Using VSAM in an MVS/XA Environment" on page 83. Detailed information on VSAM and

VSAM data sets is contained in *MVS/Extended Architecture VSAM Administration Guide*, GC26-4015.

New Device Support

Information to support the following new device has been added to Part 1, "Planning for Data Facility Product Installation," and to Appendix A, "Data Facility Product Features":

IBM 3800 Model 3 Printing Subsystem in all-points-addressable mode and Model 1 compatibility mode

New Programming Support

Information to support Version 1 Release 5 of the Resource Access Control Facility (RACF) licensed program, including a discussion of generic profiles, has been added under "Security" in Chapter 5, "Integrated Catalog Facility Catalog Installation."

Release 1.1, April 1983

Running with VSAM in 31-bit Addressing Mode

A summary of VSAM enhancements has been included in this book for planning purposes. See "Running with VSAM in 31-Bit Addressing Mode" in "Planning for VSAM Installation and Conversion."

The new VSAM enhancements include:

- VSAM record management code residing in the EPLPA
- VSAM data buffers above 16 megabytes
- Multiple LSR pools per address space (with I/O buffers above or below 16 megabytes)

Summary of Data Facility Product Features

An appendix, "Data Facility Product Features" has been added. The appendix lists new and changed features of Data Facility Product, indicates which macros and commands are affected, and states where information about each feature is documented.

Support for Global Resource Serialization

Global resource serialization (GRS) allows you to use cross-region sharing among a group of interconnected processors. A note concerning global resource serialization has been added to the section "Control Block Update Facility (CBUF)."

Contents

Figures

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 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: 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 will be 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 Version 1 (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 (5664-284), 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 for Releases 1.0, 1.1, and 1.2. (Service-updated versions of MVS/XA Data Facility Product Releases 1.0 and 1.1 will be distributed with Release 1.2.) Release 1.0 must be installed on an MVS/370 or OS/VS2 MVS base operating system that contains the appropriate level of device and programming support software for your current device configuration. Releases

1.1 and 1.2 of MVS/XA Data Facility Product may be installed concurrently with Release 1.0, or may be added to an existing MVS/XA system that has Release 1.0 of MVS/XA Data Facility Product already installed.

The OS/VS2 MVS, MVS/370, or MVS/XA base operating system, with appropriate device or programming support installed, is referred to as the **installing system** or the **generating system**. These three types of base systems, OS/VS2 MVS, MVS/370, and MVS/XA, are discussed in the following sections.

OS/VS2 MVS

In order to install MVS/XA Data Facility Product and its corequisites 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. See "Programs Replaced by Data Facility Product" on page 11 for more information on the functions of these programs.

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

If you already have MVS/370 Data Facility Product (5665-295) installed, you can convert your MVS/370 operating system to MVS/XA by installing MVS/XA Data Facility Product 2.1.0 (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. See "Programs Replaced by Data Facility Product" on page 11 for more information on the functions of these programs.

MVS/XA Data Facility Product replaces the functions provided by MVS/370 Data Facility Product. However, prior to installing MVS/XA DFP, you must 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)

If you have already installed MVS/XA Data Facility Product Release 1.0, you may add Releases 1.1 and 1.2 to your existing MVS/XA system.

Before installing Release 1.1 or 1.2, plan to use the RECEIVE and APPLY CHECK functions of the System Modification Program (SMP or SMP/E) to assure that the service level of your MVS/XA system matches the maintenance

required for that release. If the results of the APPLY CHECK indicate that your existing MVS/XA system requires the maintenance supplied in the updated version of MVS/XA Data Facility Product, you should either:

*   Apply all outstanding PTFs against MVS/XA Data Facility Product Releases 1.0 and 1.1 (a list of outstanding PTFs is available from your IBM representative); or

*   Plan to reinstall Release 1.0 when installing Releases 1.1 and 1.2. Service-updated versions of Releases 1.0 and 1.1 will be distributed with Release 1.2.

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 requires a processor complex, such as the 308x, capable of running in extended architecture mode.

MVS/XA Data Facility Product can be installed using any processor that is 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 3081, 3083, 3084, 4381, 9081, and 9083 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 following table summarizes the most commonly used I/O devices supported by the MVS/XA operating system. See *System Generation* for more information.

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

The Direct Access Storage Device Migration Aid licensed program (5668-002)
can be used to assist in migrating to the IBM 3375 or 3380 Direct Access Storage
Devices. See *DASD Migration Aid General Information* for details.

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.

See "Programs Replaced by Data Facility Product" on page 11 for a list of
devices that require Data Facility Device Support if they are to be used during
sysgen.

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 IBM 3800 Model 3 can also be used in full function mode and 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 IBM 3800 Model 3, some programming support is required in addition to Data Facility Product. See "Device Support" on page 11 for details on programming requirements.

For more information on the functions of the 3800 Model 3, see *IBM 3800 Printing Subsystem Model 3 Programmer's Guide: Compatibility.*

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

In order to sysgen MVS/XA Data Facility Product, you must have MVS/System Product 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 Release 1.0 requires at least Version 2, Release 1.2 of MVS/SP JES2 (5740-XC6), or MVS/SP JES3 (5665-291). Later releases of MVS/SP Version 2 may be used, and may also be required to support specific features or devices.

The following tables list the MVS/SP JES2 or JES3 releases 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. See "Requirements for Data Facility Product and Its Features" on page 9 for more information about software requirements for Data Facility Product features.

Device	MVS/SP JES2 or JES3 Version
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
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

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)

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

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

Assembler H Version 2 (5668-962) is required to perform the assemblies needed for system generation and maintenance of MVS/XA Data Facility Product. 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

System Modification Program Extended (SMP/E) (5668-949) or System Modification Program (SMP) Release 4.27 is required in order to install Data Facility Product. For more information on use of the System Modification Program (SMP/E or SMP), see "System Modification Program (SMP/E or SMP)" on page 17.

Environmental Recording Editing and Printing: EREP

Release 2.1 of the Environmental Recording Editing and Printing (EREP) licensed program (5752-VS2) is required for SYS1.LOGREC processing under MVS/XA. Any subsequent release of EREP will contain the same support.

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

Device Support Facilities

The MVS/XA version of Device Support Facilities (5655-257), or its equivalent, is required to initialize DASD volumes, to write IPL text, and to create indexed VTOCs for your system. Release 6 of Device Support Facilities is required to support the 3380 Direct Access Storage attached to the 3880 Storage Control Model 13; for all other DASD, 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.

See *Device Support Facilities User's Guide and Reference* for more information on Device Support Facilities.

Data Facility Data Set Services: DFDSS

Data Facility Data Set Services (5740-UT3), or its equivalent, is required to back up DASD files on tape, in a format that can be successfully restored following installation of MVS/XA Data Facility Product.

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 only restore 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. See "Programs Not Supported by Data Facility Product" on page 12 for a discussion of alternatives to previous IBM dump and restore programs.

Release 1.1 of DFDSS contains support for the IBM 3375 and 3380 Direct Access Storage devices. Release 1.2 contains support for the 3380 Direct Access Storage attached to the IBM 3880 Storage Control Model 13. Release 2.1 contains support for the 3800 Printing Subsystem Model 3.

See *DFDSS: General Information* for more information on DFDSS.

Requirements for Data Facility Product and Its Features

Data Facility Product Base

The programs listed in the following table 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 Version 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	5.0 with PTF or 6.0
Data Facility Hierarchical Storage Manager (DFHSM) Version 2	5665-329	1.0
Information Management System/Virtual Storage (IMS/VS) Version 1	5740-XX2	2.0 + compatibility feature

Licensed Program	Number	Release
Resource Access Control Facility (RACF) Version 1	5740-XXH	4.0 (for DFP 1.0 or 1.1) 5.0 (for DFP 1.2)
Resource Measurement Facility (RMF) Version 3	5665-274	1.0
DFSORT	5740-SM1	5.0

Notes:

1. *Different levels of these products may also be required to support specific devices and features under MVS/XA Data Facility Product. See "Device Support" on page 11 for a list of devices requiring additional software support. See "Data Facility Product Features" for a list of features requiring additional software support.*

2. *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, then the PTF number was not available at the time this manual was published. See the Data Facility Product program directory for detailed information on PTF numbers.*

Chapter 3, "Updating the Publications Library" on page 19, contains a list of publications related to these licensed programs.

Data Facility Product Features

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

Feature	Products Required	Program Number
Enciphering and deciphering data using access method services	Programmed Cryptographic Facility or Cryptographic Unit Support *	5740-XY5 or 5740-XY6
RACF generic profile support	RACF Version 1 Release 5	5740-XXH

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

Different levels of these programs may also be required to support specific devices under MVS/XA Data Facility Product. See "Device Support" on page 11 for a list of devices that require additional software support.

Chapter 3, "Updating the Publications Library" on page 19, contains a list of publications related to these programs.

Device Support

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

Device	Programs Required	Program Number
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, then the PTF number was not available at the time this manual was published. See the Data Facility Product program directory for detailed information on PTF numbers.

Incompatible Programs and Licensed Programs

Programs Replaced by Data Facility Product

The functions of the 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 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 just listed.

The following table indicates which levels of the replaced programs contain support for various functions, and which level of Data Facility Product replaces those functions. Any later release of Data Facility Product will also contain the appropriate support.

Program	Release	Functions Supported	MVS/XA DFP Level
Data Facility Device Support	1.0	Indexed VTOC Alternate master catalog 3375, 3380 Direct Access Storage	1.0

Program	Release	Functions Supported	MVS/XA DFP Level
	1.5	3880 Storage Control Model 11 3880 Storage Control Model 13	1.0
	1.6	3800 Printing Subsystem Model 3 in Model 1 compatibility mode	1.2
Data Facility Extended Function	1.0	Integrated catalog facility catalogs VSAM extensions	1.0
SAM-E	1.0	SAM extensions	1.0
Access Method Services Cryptographic Option	—	Enciphering and deciphering of data	1.0
Offline 3800 Utility	—	3800 Printing Subsystem Model 1 operating in offline mode	1.0

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 following table.

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) (5740-UT3) 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) (5740-UT3): restore functions only

Old Program	Function	Alternative
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) (5740-UT3): 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 program products 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 can be found in *System Generation*.

Data Facility Product

Storage Requirements

Detailed information about storage requirements for Release 1.2 will be found in the program directory for Data Facility Product.

Installation Requirements

To install MVS/XA Data Facility Product Release 1.0, you must execute a complete sysgen (GENTYPE = ALL). Assembler H Version 2, SMP Release 4.27 or SMP/E, 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 program product (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 4.27	Installing MVS/XA DFP Using SMP/E
1. Update your existing distribution libraries for Data Facility Product using the ACCEPT function of SMP.	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.	2. Execute a Stage I sysgen with GENTYPE = ALL.
3a. If you are installing MVS/XA Data Facility Product Release 1.0, 1.1, and 1.2 simultaneously, execute Stage II of a complete sysgen. New devices may be added to the system at this time. 3b. If you are installing MVS/XA Data Facility Product Release 1.2 on a system that already has Release 1.0 installed, update the system libraries using the JCLIN and APPLY functions of SMP. Once you have applied Data Facility Product, if you wish to install any new devices, you will need to execute Stage I and II of an I/O device generation (iogen) to update your I/O configuration.	3. You may skip Stage II by using the JCLIN and APPLY functions of SMP/E to update your system libraries. Once you have applied Data Facility Product, if you wish to install any new devices, you will need to execute Stage I and II of an I/O device generation (iogen) to update your I/O configuration. If you choose not to skip Stage II, new devices may be added to the system at this time.

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 program products that are not supported by the system generation process.

See *System Generation* for details on planning for and performing a complete sysgen or iogen.

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 Program Products

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) Release 1, Modification Level 0 or later (5668-949, with PTF UR90041)

- System Modification Program (SMP) Release 4, Modification Level 27 or later (with PTF UR03129)

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. See the Data Facility Product program directory for more information on using a temporary SMPMTS data set.

Device Support Facilities

The Device Support Facilities system control program requires approximately 384K bytes of virtual storage and 200K bytes of auxiliary 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 program product, if used, 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. See Appendix D of *DFDSS: User's Guide and Reference* for details on how to estimate the amount of storage a specific DFDSS operation will require.

Chapter 3. 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.

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

MVS/XA 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, or DFEF books in the right column.

MVS/XA Book	OS/VS or DFDS or DFEF Book(s)
*MVS/XA Access Method Services Logic, LY26-3889	OS/VS2 Access Method Services Logic, SY35-0010 Data Facility Extended Function: Access Method Services Logic, LY26-3888
*MVS/XA BDAM Logic, LY26-3893	OS/VS2 BDAM Logic, SY26-3831
MVS/XA Cache Device Administration, GC26-4017	None

MVS/XA Book	OS/VS or DFDS or DFEF Book(s)
MVS/XA Catalog Administration Guide, GC26-4041	*Data Facility Extended Function: Access Method Services Reference*, SC26-3967
	Data Facility Extended Function: Access Method Services Administration and Services, SC26-3966
	Planning for Enhanced VSAM Under OS/VS, GC26-3842
	OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838
	OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications, GC26-3819
	OS/VS2 MVS CVOL Processor, GC26-3864
MVS/XA Catalog Diagnosis Guide, SY26-3899	*Data Facility Extended Function: Catalog Diagnosis Guide*, SY26-3887
MVS/XA Catalog Diagnosis Reference, SY26-3897	*Data Facility Extended Function: Catalog Diagnosis Reference*, SY26-3886
MVS/XA Checkpoint/Restart Supervisor Call Logic, LY26-3890	*OS/VS2 Checkpoint/Restart Logic*, SY26-3820
MVS/XA Checkpoint/Restart User's Guide, GC26-4012	*OS/VS2 MVS Checkpoint/Restart*, GC26-3877
MVS/XA Common VTOC Access Facility Diagnosis Reference, SY26-3929	*Data Facility Device Support: Common VTOC Access Facility Diagnosis Reference*, SY26-3882
MVS/XA CVOL Processor Logic, LY26-3895	*OS/VS2 CVOL Processor Logic*, SY26-3860
MVS/XA DADSM and Common VTOC Access Facility Diagnosis Guide, SY26-3896	*Data Facility Device Support: DADSM and Common VTOC Access Facility Diagnosis Guide*, SY26-3880
MVS/XA DADSM Diagnosis Reference, SY26-3904	*Data Facility Device Support: DADSM Diagnosis Reference*, SY26-3881
MVS/XA Data Administration Guide, GC26-4013	*OS/VS2 MVS Data Management Services Guide*, GC26-3875
MVS/XA Data Administration: Macro Instruction Reference, GC26-4014	*OS/VS2 MVS Data Management Macro Instructions*, GC26-3873
MVS/XA Data Administration: Utilities, GC26-4018	*OS/VS2 MVS Utilities*, GC26-3902
MVS/XA Data Facility Product: General Information, GC26-4007	None

MVS/XA Book	OS/VS or DFDS or DFEF Book(s)
MVS/XA Data Facility Product: Master Index, GC26-4069	None
MVS/XA Data Facility Product: Planning Guide, GC26-4040	*Data Facility Extended Function: Access Method Services Administration and Services*, SC26-3966 *Data Facility Device Support: User's Guide and Reference*, SC26-3952
MVS/XA Installation: System Generation, GC26-4009	*OS/VS2 System Programming Library: System Generation Reference*, GC26-3792
MVS/XA Integrated Catalog Administration: Access Method Services Reference, GC26-4019	*OS/VS2 Access Method Services*, GC26-3841 *OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option*, SC26-3916 *Data Facility Extended Function: Access Method Services Reference*, SC26-3967
MVS/XA Integrated Catalog Administration: Access Method Services Reference Summary, GX26-3739	None
*MVS/XA ISAM Logic, LY26-3894	*OS/VS2 ISAM Logic*, SY26-3833
MVS/XA Linkage Editor and Loader User's Guide, GC26-4011	*OS/VS Linkage Editor and Loader*, GC26-3813
*MVS/XA Linkage Editor Logic, LY26-3902	*OS/VS Linkage Editor Logic*, SY26-3815
*MVS/XA Loader Logic, LY26-3901	*OS/VS Loader Logic*, SY26-3814
MVS/XA Magnetic Tape Labels and File Structure Administration, GC26-4003	*OS/VS Tape Labels*, GC26-3795
MVS/XA Media Manager Diagnosis Guide and Reference, SY26-3898	*Data Facility Device Support: OS/VS2 Media Manager Diagnosis Guide and Reference*, SY26-3884
*MVS/XA Open/Close/EOV Logic, LY26-3892	*OS/VS2 Open/Close/EOV Logic*, SY26-3827
*MVS/XA SAM Logic, LY26-3891	*OS/VS2 SAM Logic*, SY26-3832 *OS/VS2 MVS SAM-E Logic*, LY26-3855

MVS/XA Book	OS/VS or DFDS or DFEF Book(s)
MVS/XA System—Data Administration, GC26-4010	*OS/VS2 System Programming Library: Data Management*, GC26-3830 *Data Facility Device Support: User's Guide and Reference*, SC26-3952
*MVS/XA Utilities Logic, LY26-3903	*OS/VS Utilities Logic*, SY35-0005
*MVS/XA VIO Logic, LY26-3900	*OS/VS2 VIO Logic*, SY26-3834
MVS/XA VSAM Administration Guide, GC26-4015	*OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide*, GC26-3838 *OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications*, GC26-3819 *Planning for Enhanced VSAM Under OS/VS*, GC26-3842 *Data Facility Extended Function: Access Method Services Administration and Services*, SC26-3966 *OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option*, SC26-3916
MVS/XA VSAM Administration: Macro Instruction Reference, GC26-4016	*OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide*, GC26-3838 *OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications*, GC26-3819
MVS/XA VSAM Catalog Administration: Access Method Services Reference, GC26-4075	*OS/VS2 Access Method Services*, GC26-3841 *OS/VS1 and OS/VS2 MVS Access Method Services Cryptographic Option*, SC26-3916 *Data Facility Extended Function: Access Method Services Reference*, SC26-3967
*MVS/XA VSAM Logic, LY26-3907	*OS/VS2 VSAM Logic*, SY26-3825

Product-Related Publications

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

MVS/System Product JES2 or JES3

- *MVS/System Product Version 2 General Information Manual*, GC28-1118

- *MVS/Extended Architecture JES3 Introduction*, GC23-0049

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

Device Support Facilities and Data Facility Data Set Services (DFDSS)

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

- *Data Facility Data Set Services: General Information*, GC26-3947

- *Data Facility Data Set Services: User's Guide and Reference*, SC26-3949

Other Program Products

- *Advanced Communications Function for TCAM Version 2 General Information: Introduction*, GC30-3057

- *Advanced Communications Function for VTAM Version 2 General Information*, GC27-0608

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

- *OS/VS BTAM*, GC27-6980 (with BTAM/SP supplement, SC27-0604)

- *CICS/VS General Information*, GC33-0155

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

- *Data Facility Hierarchical Storage Manager General Information*, GH35-0092

- *IMS/VS Version 1 General Information Manual*, GH20-1260

- *Resource Access Control Facility (RACF): General Information Manual,* GC28-0722

- *Resource Measurement Facility (RMF) Version 3 General Information,* GC28-1115

- *DFSORT General Information,* GC33-4033

- *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

Part 2. Planning for Integrated Catalog Facility Catalogs

Chapter 4. Integrated Catalog Facility Catalog Overview

Integrated catalog facility catalogs are functional replacements for VSAM catalogs and OS CVOLs. An integrated catalog facility catalog may 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. Installation of or migration to integrated catalog facility catalogs may be implemented at a pace chosen by the installation. All existing catalogs (VSAM and OS CVOLs) continue to operate in the same way. In general, no changes are necessary to access method services job streams, user job streams that continue to use existing catalogs, or programs (except those dependent on the format of data set names in the VTOC). For additional information about integrated catalog facility catalogs, 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 1 shows the relationship of the BCS and the VVDS.
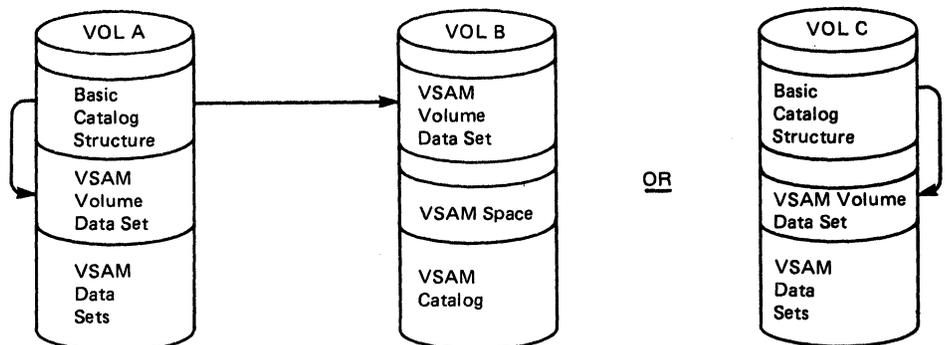
Figure 1. The Relationship of the BCS and the VVDS

The BCS

The BCS is a key-sequenced data set (KSDS) and 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 and 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, and 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 highly 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 may not be password protected, exported or imported, or have its attributes 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, or relative record

- *Alternate index*, creates an alternate index entry, and an index entry to define the alternate index

- *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 2 on page 30 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.
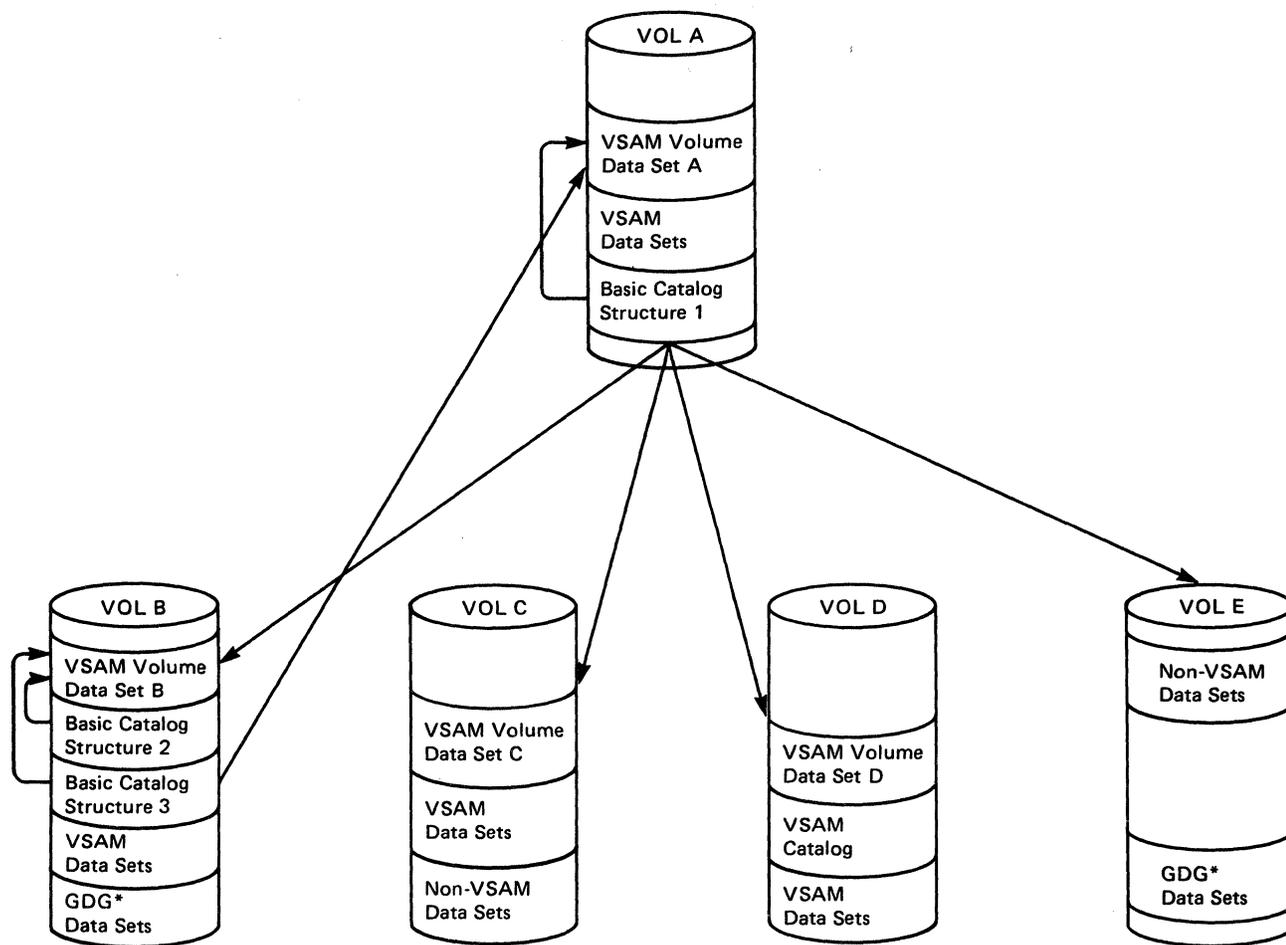
	has cataloged entries on:
Basic Catalog Structure 1	VOL A, VOL B, VOL C, VOL D, VOL E
Basic Catalog Structure 2	VOL B
Basic Catalog Structure 3	VOL A, VOL B

*Generation Data Group Data Sets

Figure 2. Integrated Catalog Facility Catalog Relationships

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 for the duration of 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*.

Chapter 5. 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 a priority sequence.

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 is not exhaustive, but 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) licensed program (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 new data set names that occur in the integrated catalog facility catalog environment.

The use of JOBCAT and STEPCAT statements should be restricted to those jobs that require them, for example, those jobs which 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*.)

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

Figure 3 on page 36 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 by either using mount commands or 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, if 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.

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 to determine appropriate actions 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 previously discussed 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.

- To conserve common service area (CSA) space, limit the number of user catalogs.
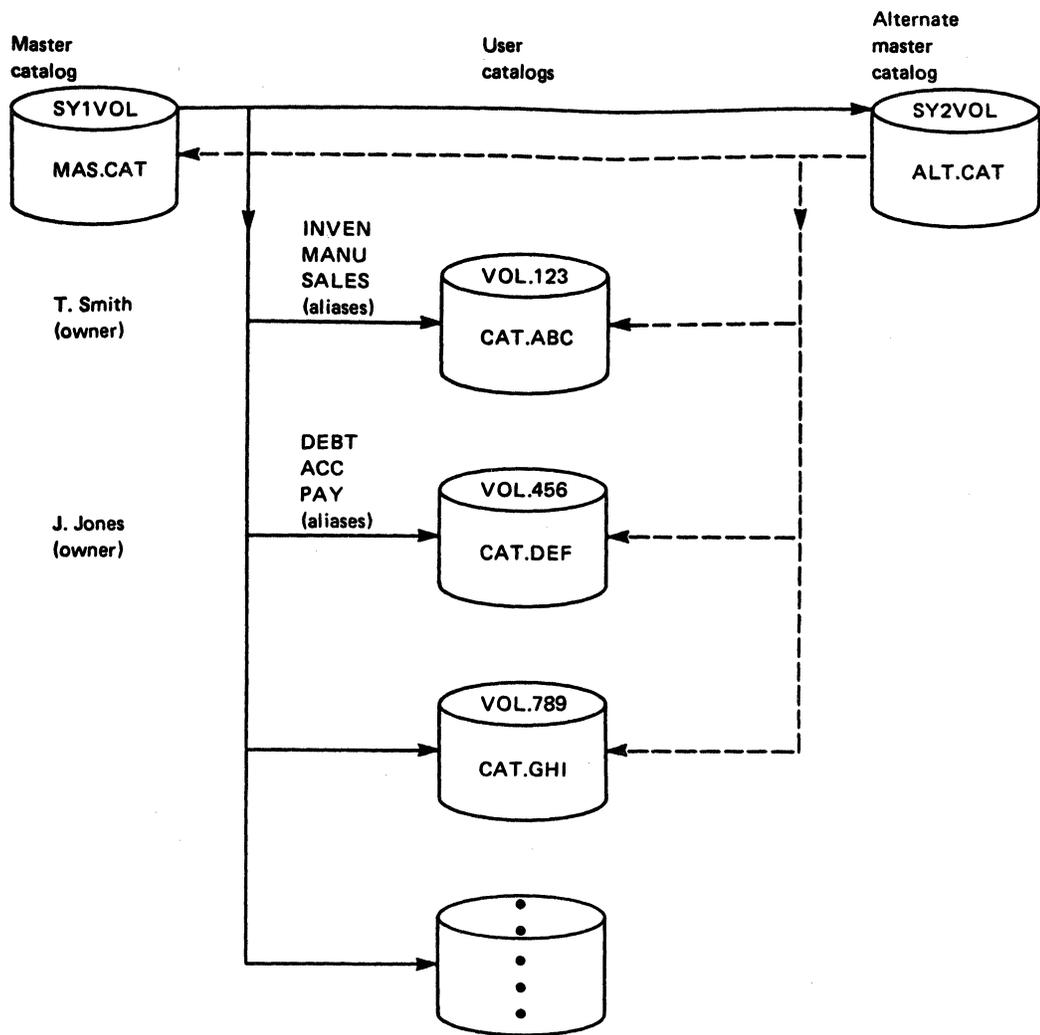
Figure 3.   Catalog Configuration Diagram

- When selecting the catalog unit, consider the contention caused by the amount of use 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, also 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 in tracks to provide consistency with the control interval size specifications.

- Let BUFNI and BUFND default to STRNO + 1. These numbers should normally be adequate for performance and still not require excessive 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 53).

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. Like 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), 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 57.)

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) may be backed up by a volume dump or by EXPORT. The frequency of backup is determined by the amount of catalog activity and the acceptable recovery time.

- If the backup is a volume dump, the entire volume must be restored because IBM utilities do not support the recovery of VSAM data sets from a volume dump. If the backup is an exported copy, it may be restored by IMPORT.

- The VVDS reflects the current status of VSAM data sets, and you may not use EXPORT or REPRO.

- The VVDS may be backed up and restored as part of a total volume operation.

- Like the VTOC, if the VVDS is unusable, it may 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 may be merged into another integrated catalog facility catalog using REPRO MERGE. Entries from the source catalog may 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.

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

Integrated catalog facility catalogs have control blocks and buffers in the CSA that require virtual storage space. Each open integrated catalog facility catalog with the minimum number of strings (2) requires approximately 8K bytes (K equals 1024) of storage for its control blocks. Each additional string requires an additional 1.2K bytes of storage.

In addition to the storage for control blocks, each open integrated catalog facility catalog also requires some buffer space, which is determined by user selection of device type, data control interval size, string number, and index control interval size. For an integrated catalog facility catalog with 2 strings, there will be a minimum of 6 buffers.

For an open integrated catalog facility catalog with a string number of 2, and a control interval size of 1K for both data and index control intervals, the total CSA space requirement is 14K (8K for control blocks and 6K for buffers). For a catalog with a string number of 7, data CI size of 2K, and index CI size of 4K, the total CSA space requirement is 56K (28K for control blocks and 28K for buffers).

For more information on determining buffer space, see DEFINE USERCATALOG in *Access Method Services Reference*.

In MVS/XA, the integrated catalog facility catalog support increases the storage requirement in the pageable link pack area by approximately 219K bytes. This decreases the maximum virtual storage available to the private area of each MVS/XA address space.

MVS/XA supports both VSAM and integrated catalog facility catalogs. After all VSAM catalogs have been converted to integrated catalog facility catalogs (the master catalog and all user catalogs), and there is no further requirement for the system to support any VSAM catalogs, the CSECTs which provide support only for VSAM catalogs may be deleted. This results in a decrease of approximately 170K bytes in the size of the pageable link pack area.

In summary, the CSA requirement for integrated catalog facility catalogs depends on the number of catalogs open at any given time, and on the data sets cataloged in each of these catalogs. If a large number of OS CVOLs are involved, a one-for-one conversion may result in excessive CSA requirements.

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 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 52. 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, for example, TSO functions.

If the master catalog contains only the required entries, it should be very static. If this is true, the master catalog may be a candidate for early conversion to relieve the volume ownership of the system packs.

When converting catalogs that contain VSAM data sets, recovery procedures require the use of an additional set of commands, and require more planning and testing. This is particularly true when such catalogs support data base applications. It may take somewhat longer to replace these catalogs with production integrated catalog facility catalogs.

When scheduling the conversion of existing catalogs, it must be remembered that all systems that need access to the integrated catalog facility catalogs must have the integrated catalog facility catalog support installed. 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 6. Integrated Catalog Facility Catalog Conversion

The conversion of a VSAM catalog or OS CVOL to an integrated catalog facility catalog consists of several stages:

1. Choosing a conversion technique
2. Checking the CNVTCAT prerequisites
3. Preparing for CNVTCAT processing
4. Backing up VSAM data sets, data spaces, and the catalog
5. Setting up the VVDS
6. Setting up the BCS
7. Converting the VSAM catalog or OS CVOL
8. Verifying the conversion has been successful
9. Revising access method services procedures
10. Examining 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 40.)

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

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 38.)

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 suballocateo VSAM data sets) in the VTOC of each volume owned by the converted catalog to allow for the new DSCBs which 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 preceding 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.

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 licensed program. If there is a VSAM catalog on the volume which 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 the VSAM source catalog aliases and 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.

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:

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 following table 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	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, then the control area size is 1 track

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

- One cylinder or greater, then 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:

Device Type	Number of Tracks
3330	3
3340	6
3350	3
3375	2
3380	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.

To optimize the performance of the BCS, it is useful to keep the number of levels in the index to a minimum. If no more than two index levels are to be used, a selected index control interval size limits the capacity of the data set.

Too large a control area size extends the index beyond two levels. The following table shows the maximum control area sizes to maintain a two level index for two different index control interval sizes. Because a second level index record can address up to 121 data control areas, the total data capacity from these allocations is also given:

Device Type	Index Control Interval Size	Data Control Area (Track)	Maximum Data Capacity (Megabytes)
3330	1024	11	14.6
	512	5	3.2
3350	1024	8	14.5
	512	3	2.6
3380	1024	3	11.2
	512	1	1.8

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 index space allocation 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 to select 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 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 56.)

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. The UNIQUE and SUBALLOCATION
parameters of the DEFINE commands will be ignored. All objects will be defined
UNIQUE. 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 CAMLIST

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 also apply to converting a master catalog. There are, however, additional items that need attention:

- An active master catalog currently in use by the system cannot be converted.
- The SYS1.NUCLEUS catalog member must be updated or a new one added.

A temporary catalog can be selected for use as a master catalog during the conversion by use of the alternate master catalog facility. The temporary 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 of the minimum sizes necessary to run the conversion, and should be independent of the catalog being converted.

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

For example, one procedure is to generate a fully operational alternate master catalog for the system. The 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 the alternate master catalog generation procedure, see *Catalog Administration Guide*.

The pointer in SYS1.NUCLEUS, member SYSCATnn, is used during nucleus initialization to find the 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.

When large numbers of OS CVOLs are converted, or when there is a constraint on CSA space, it may be advisable to minimize the number of integrated catalog facility catalogs to avoid too large a CSA requirement. For details on the CSA space requirements for an integrated catalog facility catalog, see "Virtual Storage" on page 44.

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.

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

Part 3. Planning for the Indexed VTOC

Chapter 7. 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 by the hardware of the VTOC, which make the device and channel unavailable, and manages free space information so the number of I/O operations to obtain or release space on the volume is reduced.

The VTOC index is optional. It can be built over an existing VTOC and "switched off," with or without deletion of the index. This allows DADSM requests to be processed 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 'xxxxxxxx' 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 4. For a complete description of the parts of an index, see *System — Data Administration*.

```
            VTOC                                VTOC Index
 +-------------------------------+        +--------------------------+
 |        Format-4 DSCB          |    +-->|                          |
 +-------------------------------+    |   +--------------------------+
 |        Format-5 DSCB          |    |   |                          |
 +-------------------------------+    |   +--------------------------+
 |                               |    |   |                          |
 |                               |    |   +--------------------------+
 |                               |    |   |                          |
 +-------------------------------+    |   +--------------------------+
 | Format-1 DSCB for the VTOC    |----+   |                          |
 | Index SYS1.VTOCIX.xxxxxxxx    |        +--------------------------+
 +-------------------------------+        |                          |
 |                               |        +--------------------------+
 |     Other Format-1 DSCBs      |        |            .             |
 |                               |        |            .             |
 +-------------------------------+        |            .             |
                                          +--------------------------+
```
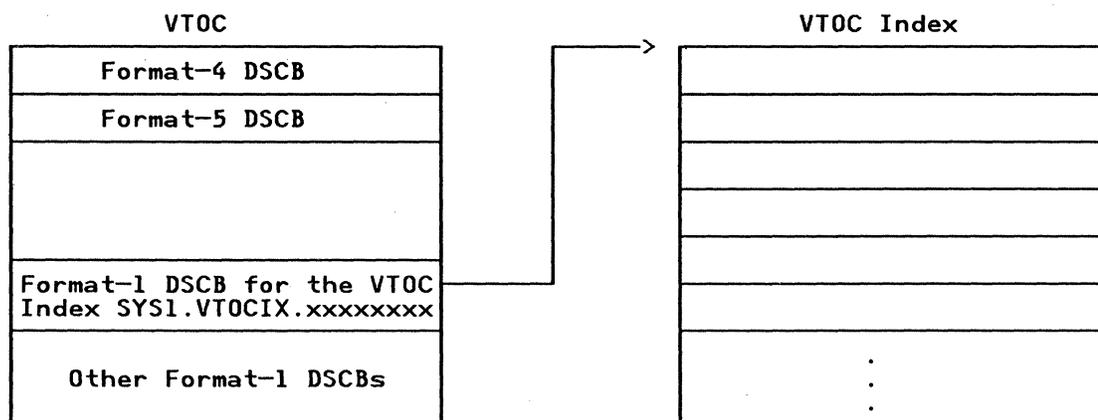
Figure 4. Relationship of a VTOC to Its Index

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. However, to assist user programs dependent upon the contents of this field to read the VTOC without going through DADSM, the field is initialized at the time the VTOC index is activated to point to the last DSCB slot in the VTOC extent.

- 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 (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.

- 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.

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, DADSM and CVAF Diagnosis Guide*, and *DADSM 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

Chapter 8. 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 of channel and device utilization and allocates a new data set, renames a data set, and obtains a data set name very efficiently. An indexed VTOC also provides significant improvements for allocating an existing data set, deleting a data set, OPEN for input, and OBTAIN by TTR.

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 which 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 alphabetical 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 then 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 VTOC is converted into 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. Use a DD statement. (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

Program compatibility is discussed in the migration section above. 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 which have 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 once 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 may 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, since 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 just like 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 which 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. These are described below as they affect data set and volume types of backup.

Data Set Backup

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

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 further 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 OS/VS2 MVS operating system and other previous licensed programs. 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
3480 Magnetic Tape Subsystem	Description of programming support	—	*Data Facility Product: General Information*
(3420 compatibility)	Specifying 3420 compatibility mode	IODEVICE macro UNIT = 3420C parameter	*System Generation*
(3480 full function)	Specifying the device during system generation	IODEVICE macro UNIT = 3480 parameter	*System Generation*
	Using 18-track tapes	—	*Magnetic Tape Labels and File Structure*
	Controlling message display	MSGDISP macro	*Data Administration Guide* *Data Administration: Macro Instruction Reference* *System — Data Administration*
	Controlling buffering	SYNCDEV macro	*Data Administration Guide* *Data Administration: Macro Instruction Reference*

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
	Using fast positioning during processing	NOTE macro POINT macro	*Data Administration Guide* *Data Administration: Macro Instruction Reference* *System — Data Administration*
3800 Model 3 Printing Subsystem	Description of programming support	—	*IBM 3800 Printing Subsystem Programmer's Guide*
	Specifying the device during system generation Print Services Facility (PSF) support	IODEVICE macro UNIT = 3800 MODEL = 3 parameters	*System Generation*
	Specifying the device for IEBIMAGE	OPTIONS statement DEVICE = 3800M3 parameter	IEBIMAGE in *Utilities*
	Specifying UCS and FCB images	SETPRT macro	*Data Administration Guide* *Data Administration: Macro Instruction Reference* *Checkpoint/Restart*
3880 Storage Control Model 11	Description of programming support	—	*Data Facility Product: General Information*
	Error messages	—	*System Messages*
	Specifying the device during system generation	IODEVICE macro UNIT = 3350P parameter	*System Generation*
3880 Storage Control Model 13	Description of programming support	—	*Data Facility Product: General Information*
	Error messages	—	*System Messages*
3880 Storage Control Model 21	Description of programming support	—	*Data Facility Product: General Information*
	Error messages	—	*System Messages*
	Specifying the device during system generation	IODEVICE macro UNIT = 3351P parameter	*System Generation*
3880 Storage Control Model 23	Description of programming support	—	*Data Facility Product: General Information*
	Error messages	—	*System Messages*

New Programming Support

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
ISO/ANSI/FIPS tape labels and file structure	Description of programming support	—	*Magnetic Tape Labels and File Structure*
	Including ISO/ANSI/FIPS support in the operating system	CTRLPROG macro ASCII = INCLUDE parameter	*System Generation*
	Specifying an access security code for ISO/ANSI/FIPS tapes in IEHINITT	INITT statement ACCESS keyword	*Utilities*
	Specifying an access security code for ISO/ANSI/FIPS tapes in JCL	ACCODE keyword	*JCL*
	Description of data management support	—	*Data Administration Guide*
	Specifying data management support	DCB macro	*Data Administration: Macro Instruction Reference*
Global resource serialization (GRS)	Description of GRS	RNAME feature	*Global Resource Serialization*
	Specifying cross-region sharing of VSAM data sets among interconnected processors	SHAREOPTION parameter	*Access Method Services Reference* *VSAM Administration Guide*
Linkage editor enhancements	Specifying 31-bit addressing mode	AMODE, RMODE parameters	*Linkage Editor and Loader*
CHKPT macro enhancements	Specifying additional addresses for a checkpoint	CHKPT macro: DDNADDR = ddnaddr IDADDR = checkid addr IDLNG = checkid length IDLNG = S MF = (S,addr) MF = (L,name) MF = (E,addr1,addr2)	*Checkpoint/Restart*

Supported Feature	Task or Topic Covered	New or Modified Macros/Commands	Documented in These Books
IEBCOPY enhancements	Altering, copying, and reblocking load modules	ALTERMOD statement COPYMOD statement	*Utilities* *System Messages*
IEHMOVE enhancements	Multiple BSAM buffering	REGION parameter (JCL)	*Utilities* *System Messages* *JCL*
31-bit VSAM	Specifying 31-bit addressing mode with VSAM	BLDVRP, DLVRP, ACB macros GENCB, MODCB, SHOWCB, TESTCB macros	*VSAM Administration Guide* *VSAM Administration: Macro Instruction Reference*
RACF generic profile support	Description of RACF Release 5 and generic profiles	—	*RACF General Information Manual*
	Effect of RACF and generic profiles on catalog administration	—	*Access Method Services Reference* *Catalog Administration Guide* *VSAM Administration Guide*

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 Resides in the ELPA

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

VSAM code in the ELPA is accessed by applications residing below 16 megabytes through the aid of an interface module. No changes to user programs are required.

VSAM Data Buffers Can Be Obtained above 16 Megabytes

User programs that run in 31-bit addressing mode can create and access VSAM data buffers from any part of virtual storage, including the area above 16 megabytes.

To use the extended VSAM buffer support, users must specify 31-bit buffering in the ACB at OPEN time. The ACB must remain below 16 megabytes, although the RPL may reside above or below that point.

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 LSR (local shared resources) pools in an address space. Up to 16 LSR pools can be built for each address space, and I/O buffers may reside above or below 16 megabytes.

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. (You must code the AMODE = 31 parameter on the ACB and BLDVRP macros.)

• User programs requesting data buffers above 16 megabytes must run in 31-bit addressing mode.

- 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 run 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.

- You must observe the following restrictions regarding macro instructions and parameter lists that must reside below 16 megabytes:

Macro Instruction	Residence When Issued	Residence of PARM List	Other Residence Restrictions
ACB	below 16MB	N/A	below 16MB
OPEN, CLOSE	any	below 16MB	ACB below 16MB
BLDVRP, DLVRP	any	below 16MB	
RPL, EXLST	any	any	ACB below 16MB
GENCB, MODCB, SHOWCB, TESTCB	any	any	ACB below 16MB
ACQRANGE, CNVTAD, GET, GETIX, PUT, PUTIX, CHECK, ENDREQ, ERASE, MNTACQ, MRKBFR, POINT, SCHBFR, VERIFY, WRTBFR	any	any	ACB below 16Mb

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 Dictionary of Computing, SC20-1699*.

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.

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.)

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

controlling integrated catalog facility catalog
performance options   33
rebuilding alias entries after conversion to integrated
catalog facility   56, 61
recovering the BCS   39, 41
reducing maximum record size for the BCS   55
reorganizing the BCS   42
space required when building entries in the BCS   56
DELETE command
recovering the BCS   41
determining which VTOCs should be indexed   70
Device Support Facilities   8, 13
devices
compatible with Data Facility Product   5
DFDSS (Data Facility Data Set Services)   9, 13, 52
DIAGNOSE command
locating discrepancies between the BCS and VSAM
data sets   41
verifying conversion to integrated catalog facility   58
VSAM data set recovery time   35
direct access storage device (DASD)
compatible with Data Facility Product   5
DOS bit
off   72, 73
on   66

E

entry
in a catalog   29
EREP   8
EXPORT command
backing up data and associated catalog entries   52
backing up the integrated catalog facility
catalog   40
moving VSAM data sets when volume is full   51
removing user catalog connector   56
reorganizing the BCS   42, 56
verifying conversion to integrated catalog facility   57
EXPORTRA command
not supported by integrated catalog facility   59
releasing space for conversion to integrated catalog
facility   51

F

free space
fragmented   70
information in VTOC index   65, 70
preparing for CNVTCAT processing   51
reorganizing the BCS after conversion to integrated
catalog facility   57
full volume   51
full VTOC   51

G

generation data group   29
glossary   85

I

I/O resources   45
IBM licensed programs
compatible with Data Facility Product   8
incompatible programs   8
programs required for new features   8
IEHLIST program
detecting invalidation of VTOC index   73
used to manage OS CVOL   61
verifying conversion to integrated catalog facility   57
IEHMOVE program
used to manage OS CVOL   61
IEHPROGM program
meets requirements of authorized program
facility   76
protecting the VTOC index   75
restricted to selected users   40
used to manage OS CVOL   61
IMASPZAP program   75
IMPORT command
backing up the integrated catalog facility
catalog   42
moving VSAM data sets when volume is full   51
providing free space   57
reorganizing the BCS   42, 56
verifying conversion to integrated catalog facility   57
IMPORTRA command
not supported by integrated catalog facility   59
releasing space for conversion to integrated catalog
facility   51
incompatible programs   8
with Data Facility Product   8
index bit   66, 72
index bit reset feature   74
INIT command   71
initializing DASD volumes   18
installations with VSAM catalogs   59
installing
Data Facility Product   3
integrated catalog facility
See integrated catalog facility catalog
integrated catalog facility catalog
See also BCS, VVDS
catalog
backup and recovery   58
configuration   34
constructing   37
conventions   32

Numerics

MVS/XA Data Facility Product:
Planning Guide
GC26-4040-3

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:

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.)

GC26-4040-3

Reader's Comment Form

BUSINESS REPLY MAIL
FIRST CLASS   PERMIT NO. 40   ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
Programming Publishing
555 Bailey Avenue
P.O. Box 49023
San Jose, Ca. 95161-9023

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

IBM®

IBM
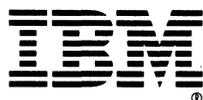
MVS/Extended Architecture
Data Facility Product:
Planning Guide

S370-34

GC26-4040-03