

GC26-3841-3
File No. S370-30

Systems

OS/VS2 Access Method Services

Release 3.8

IBM

This publication was produced using the
IBM Document Composition Facility
(program number 5748-XX9)
and the master was printed on the
IBM 3800 Printing Subsystem.

Fourth Edition (October 1980)

This is a major revision of, and makes obsolete, GC26-3841-2.

This edition applies to Release 1.0 of Data Facility/Device Support, Program Product 5740-AM7, as well as to Release 3.8 of OS/VS2 MVS, and to any subsequent releases until otherwise indicated in new editions or technical newsletters.

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

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

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Publications are not stocked at the address given below; requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

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

© Copyright International Business Machines Corporation 1975,
1976, 1978, 1980

ABOUT THIS BOOK

This book describes the use of Access Method Services commands, a group of utility functions vital to Virtual Storage Access Method (VSAM). This publication provides VSAM information required to use Access Method Services to establish and maintain data sets. For information on VSAM data set and catalog format and structure, see Planning for Enhanced VSAM Under OS/VS, GC26-3842. For information on the use of VSAM macro instructions, VSAM optimization options, and various VSAM algorithms affecting performance, see OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838.

Readers of this book are presumed to have a background in programming.

This book has the following major divisions:

- "Guide to Access Method Services," which lists functions and the Access Method Services commands used to perform them.
- "Introduction," which provides an overview of Access Method Services, including general language considerations and the use of Time Sharing Option (TSO) and the 3850 Mass Storage System with VSAM and Access Method Services. Everyone who intends to use Access Method Services should first read this chapter.
- "Data Security and Protection," which describes the recovery features of VSAM that are specified as options in Access Method Services commands.
- "Creating and Cataloging Objects," which describes the use of the DEFINE command to define catalogs, data spaces, VSAM data sets, alternate indexes, paths, and page spaces, and to catalog generation data groups, aliases and nonVSAM data sets.
- "Building an Alternate Index," which describes the use of the BLDINDEX command to build an alternate index. An alternate index allows you to access a data set with a new set of key values.
- "Modifying Catalog Information," which describes the use of the ALTER command to modify attributes in catalog entries.
- "Displaying Catalog Information," which describes the use of the LISTCAT command to list catalog entries.
- "Deleting Catalog Entries," which describes the use of the DELETE command to delete catalog entries.
- "Moving Entries," which describes the use of the IMPORT and EXPORT commands to create a backup copy of a data set or to move a data set or user catalog between systems.
- "Converting an OS Catalog's Entries to VSAM Catalog Entries," which describes the use of the CNVTCAT command to convert OS catalog entries to VSAM catalog entries and merge them into a VSAM catalog.
- "Restoring a Cluster's End-of-File Values," which describes the use of the VERIFY command to ensure that the true end-of-file is reflected in the catalog.
- "Restoring Catalog Entries After System Failure," which describes the use of the LISTCRA, EXPORTRA, IMPORTRA, and RESETCAT commands to retrieve the contents of VSAM data sets after a VSAM catalog is damaged.

- "Copying and Printing," which describes the use of the REPRO command to copy and to reorganize data, and to convert data from indexed-sequential or sequential organization to VSAM organization and from VSAM organization to sequential organization, and to copy catalogs; and the use of the PRINT command to print VSAM and nonVSAM data sets.
- "Listing Tape Volumes Mounted at Checkpoint," which describes the use of the CHKLIST command to list tape data sets that were open during a checkpoint.
- "Command Format," which lists the Access Method Services commands and their format, describes the required and optional parameters, and specifies their abbreviations. There are examples with each command and a description of how the commands and their parameters were used.
- "Controlling Command Execution," which describes the use of the IF, SET, and PARM commands to control command execution and to specify diagnostic aids and printed-output options.
- "Appendix A: Examples of Jobs Using Access Method Services Commands," which describes how two or more Access Method Services commands can be used together to accomplish a task.
- "Appendix B: Interpreting LISTCAT Output Listings," which provides information on the structure of LISTCAT output and shows sample output.
- "Appendix C: JCL DD Parameters to Take Care With," which describes JCL parameters that have either no effect in a VSAM environment or that have a negative effect.
- "Appendix D: Interpreting LISTCRA Output Listings," which describes the format and uses of LISTCRA output.
- "Appendix E: Sample Output from CHKLIST," which shows a CHKLIST listing and explains its contents.
- "Appendix F: Command Reference Summary" which shows the format of each Access Method Services command, parameter abbreviations, and defaults. The commands are grouped together to allow you to remove these pages and use them as a quick reference guide.
- "Appendix G: Invoking Access Method Services from a Problem Program," which describes how Access Method Services may be dynamically invoked from a user's problem program.
- "Appendix H: Changing the Attributes of the VSAM Master Catalog," which describes how to apply new attributes to the VSAM master catalog; for example, changing it from nonrecoverable to recoverable.
- "Glossary," which defines terms relevant to Access Method Services and VSAM.
- "Index," which is a subject index to the book.

REQUIRED PUBLICATIONS

The reader should be familiar with information presented in the following publications:

- OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide, GC26-3838, which provides information on VSAM optimization options and various VSAM algorithms affecting performance, macro instructions used to process VSAM data sets, and how to use an ISAM processing program to process a VSAM data set or an indexed-sequential data set that has been converted to VSAM format.

- Planning for Enhanced VSAM Under OS/VS, GC26-3842, which includes information on the structure and contents of VSAM catalogs and data sets and also information on protecting VSAM catalogs and data sets.
- OS/VS2 MVS Utilities, GC26-3902, which describes the utility programs available for use with nonVSAM data sets on OS/VS2 systems.
- OS/VS2 MVS Data Management Services Guide, GC26-3875, which presents basic concepts such as access method, direct-access storage, and the distinction between data-set organization and data-set processing.
- OS/VS2 JCL, GC28-0692, which describes the JCL parameters referred to in this publication.
- OS/VS Message Library: VS2 System Messages, GC38-1002, which provides a complete listing of the messages issued by Access Method Services in a VS2 system.
- OS/VS2 System Programming Library: System Generation Reference, GC26-3792, which provides information necessary to install an MVS system control program.

RELATED PUBLICATIONS

The following publications contain information that is related to this publication:

- OS/VS Virtual Storage Access Method (VSAM): Options for Advanced Applications, GC26-3819, which provides information about advanced applications of VSAM which the reader doesn't need to know about to make normal use of VSAM. The topics covered include: gaining access to control intervals; I/O buffering; constructing parameter lists for the macros that generate, modify, and examine control blocks at execution; and processing an index as data.
- OS/VS2 Checkpoint/Restart, GC26-3877, which explains how to take checkpoints and to restart processing.
- Introduction to the IBM 3850 Mass Storage System (MSS), GA32-0028, which introduces the storage devices and programming of the Mass Storage System.
- OS/VS Mass Storage System (MSS) Services: Reference Information, GC35-0017, which describes the Mass Storage System Access Method Services commands for the Mass Storage System.
- OS/VS2 System Programming Library: Data Management, GC26-3830, which describes how to use the PROTECT macro.
- OS/VS2 System Programming Library: Supervisor, GC28-0628, which describes the authorized program facility.
- OS/VS2 TSO Command Language Reference, GC28-0646, which describes TSO commands.
- OS/VS2 TSO Terminal User's Guide, GC28-0645, which describes commands available to the TSO user.
- OS/VS2 System Programming Library: TSO, GC28-0629, which describes the TSO facilities that can be influenced by the system programmer.
- OS/VS2 MVS TSO Command Package: User's Guide and Reference, GC28-0748, which describes the TSO facilities that can be influenced by the system programmer.

- OS/VS2 Using OS Catalog Management with the Master Catalog: CVOL Processor, GC35-0010, which describes the use of control volumes.
- OS/VS2 Virtual Storage Access Method (VSAM) Logic, SY26-3825, which describes the functional operation of VSAM (Record Management, Open, Close, End-of-Volume, and Control Block Manipulation commands).
- OS/VS2 Catalog Management Logic, SY26-3826, which describes the functional operation of VSAM Catalog Management.
- OS/VS2 Access Method Services Logic, SY35-0010, which describes the functional operation of Access Method Services Programs.
- OS/VS2 MVS Resource Access Control Facility (RACF) General Information Manual, GC28-0722, which contains initial planning information for RACF.
- OS/VS2 Resource Access Control Facility (RACF) Command Language Reference, SC28-0733, which describes the syntax and function of RACF commands.
- OS/VS2 System Programming Library: Debugging Handbook Volume 1, GC28-0708, and Volume 2, GC28-0709, which describes debugging procedures. (Both manuals may be ordered using GBOF-8211.)

NOTATION CONVENTIONS

A uniform system of notation describes the format of Access Method Services commands. This notation is not part of the language; it simply provides a basis for describing the structure of the commands.

The command-format illustrations in this book use the following conventions:

- Brackets [] indicate an optional parameter.
- Braces { } indicate a choice of entry; unless a default is indicated, you must choose one of the entries.
- Items separated by a vertical bar represent alternative items. No more than one of the items may be selected.
- An ellipsis ... indicates that multiple entries of the type immediately preceding the ellipsis are allowed.
- Other punctuation (parentheses, commas, spaces, etc.) must be entered as shown. A space is indicated by b.
- **Boldface** type indicates the exact characters to be entered. Such items must be entered exactly as illustrated (in upper case, except in TSO).
- Lowercase type in the command format specifies fields to be supplied by the user.
- Underscored type indicates a default option. If the parameter is omitted, the underscored value is assumed.
- A b in the command format indicates a blank must be present before the next parameter.

SUMMARY OF AMENDMENTS

OS/VS2 MVS DATA FACILITY/DEVICE SUPPORT RELEASE 1.0 PROGRAM NUMBER 5740-AM7

The system residence volume serial number and/or device type can now be changed without having to recatalog all the non-VSAM data sets on the volume. See the changes for DEVICETYPES and VOLUMES parameters in the "Command Format" chapter, "DEFINE NONVSAM."

The size of an allocation unit for the IBM 3380 and IBM 3375 is included in "Creating and Cataloging Objects," section "How Space Is Assigned to a Catalog."

Editorial changes have been made throughout this book; service changes are noted by a revision bar in the left margin.

JULY 1978

Command parameters have been arranged in alphabetic order.

Appendix F has been completely revised. It is now a reference summary showing the syntax of each command, parameter abbreviations, and defaults.

Numerous editorial and service changes have been made throughout the book.

OS/VS2 MVS SUPERVISOR PERFORMANCE #2 (VS2.03.807)

AUXILIARY STORAGE MANAGER (ASM) REDESIGN

SWAP space data sets can now be defined and preformatted. Also, a new integrity attribute has been added to both SWAP and PAGE spaces.

RESOURCE ACCESS CONTROL FACILITY (RACF)

Changes are included to support protection of VSAM data sets through RACF.

OS/VS2 MVS DATA MANAGEMENT (VS2.03.808)

EXTENDED CVOL SUPPORT

In OS/VS2 MVS Release 2, the VSAM catalog replaced the system master catalog and support of OS CVOLs was limited to a subset of what had been provided in OS/VS2 SVS Release 1. Extended CVOL support provides a CVOL function that is equivalent to SVS while retaining the VSAM master catalog as the only system master catalog. This support required that changes be made to the section "VSAM's Use of Catalogs" and the chapter "Converting an OS Catalog's Entries to VSAM Catalog Entries."

RESETTING A CATALOG

A new Access Method Services command, RESETCAT, has been added to support recoverable catalogs. The new command will allow you to reset your VSAM catalog to the level of its owned volumes. That is, if the catalog or any of its volumes become damaged, RESETCAT can be used to synchronize a catalog to a restored (not current) level version of a volume. Or, if the catalog has to be restored or reloaded, its entries can be changed to reflect the current level of its volumes with RESETCAT. The new sections "Resetting a Catalog" and "RESETCAT" have been added to describe and define parameters for this command.

RELEASE 3.7

LISTCAT Format

The format of LISTCAT output has been modified to improve readability.

Miscellaneous Topics

Two new appendixes have been added. They are: Appendix G: Invoking Access Method Services from a Problem Program, and Appendix H: Changing the Attributes of the VSAM Master Catalog.

All Access Method Services command formats, parameter descriptions, and examples have been moved to a new chapter: Command Format. In the new chapter, commands are arranged in alphabetical order.

Information relating to such topics as VSAM Data Structure, How Data is Physically Stored, Types of VSAM Data Sets, VSAM Catalog Structure, Performance Optimization, and Data Integrity has been removed from this book. The information removed can be found in other VSAM publications.

INDEPENDENT COMPONENT RELEASE

Following is a summary of major changes to Access Method Services for the release of enhanced VSAM as an independent component under OS/VS2.

Alternate Index

The alternate index allows you to access data records using a key field other than the key field established when the data set was defined (for key-sequenced VSAM data sets), and to access the data records of an entry-sequenced data set with a key field.

Reusable Data Set

The reusable data set is an attribute that can be applied to VSAM entry sequenced, key sequenced, relative-record and alternate index data sets. This attribute allows the data set to be used for temporary storage of data. Whenever a reusable VSAM data set is opened, its high-used RBA can be reset to zero.

Spanned Records

Data records in VSAM data sets can be larger than one control interval, and are contained in two or more control intervals.

Relative-Record Data Set

The relative-record data set is a VSAM data set whose records are fixed-length and are sequenced according to record number.

VSAM Volume Cleanup

The VSAM volume cleanup process allows you to erase VSAM data from a volume's VTOC and reset the VTOC's format-4 DSCB to show system ownership of the volume. The process is to be used when a system or device failure occurs that damages the volume's information in its VSAM catalog (that is, the VSAM catalog that owns the volume).

VSAM Catalog Cleanup

The VSAM catalog cleanup process allows you to delete a VSAM catalog without first deleting each of its cataloged entries. You can use this function to remove an unwanted catalog from your system. (Certain safeguards prevent indiscriminate use of the function.)

Record Replacement (During REPRO)

You can merge two data sets so that, when duplicate records are encountered, the older record (that is, the one in the target data set) is replaced with the newer record (the one in the source data set).

Catalog Recovery

Three new Access Method Services commands support recoverable catalogs: LISTCRA, EXPORTRA, and IMPORTRA. The catalog recovery function allows you to repair damaged entries in a VSAM catalog when a system failure occurs.

Catalog Recovery Area

Copies of each catalog record are maintained (on the volume containing the catalog entry's data records) in an area called the catalog recovery area. The catalog recovery area is built and maintained only for recoverable catalogs, and is used to recover from damaged catalog entries.

Exception Exit Routine

You can write an I/O error handling routine, called the exception exit routine, that is tailored for a cluster's or alternate index's data or index component, and which is called before the user program's SYNAD routine gets control.

Miscellaneous Topics

The KEYS and RECORDSIZE parameters of the DEFINE CLUSTER command have been modified so that a default value exists for them. The ALTER command's parameter set now includes the KEYS and RECORDSIZE parameters.

You can import an exported object into a predefined object.

You can specify blocksize for exported and imported objects.

The dname subparameter of the CATALOG parameter has been omitted from the syntax of those commands that no longer require its specification. If you specify dname, it will be ignored.

CONTENTS

Guide to Access Method Services	1
Introduction	4
Access Method Services Commands	4
Functional Commands	5
Modal Commands	6
Language Considerations	6
Parameter Set	6
Continuing Commands	8
Continuation Cautions	8
Terminator	9
JCL and Dynamic Allocation	9
Output Data Sets	11
Invoking Access Method Services	12
As a Job or Jobstep	12
From a TSO Terminal	13
From a User's Program	14
Authorized Program Facility	14
VSAM Data Sets	15
VSAM Volume Ownership	16
VSAM's Use of Catalogs	19
The Master Catalog	20
Transporting User Catalogs	21
JOB CAT and STEPCAT DD Statements	21
OS Control Volumes (CVOLs)	23
Time Sharing Option (TSO)	24
Mass Storage System (MSS)	25
Data Security and Protection	27
Data-Set Security	27
Passwords to Authorize Access	27
Operator Prompting Code	30
Attempts to Supply a Password	30
Passwords for NonVSAM Data Sets	30
User-Security-Verification Routine	31
Resource Access Control Facility (RACF) Protection	31
Protecting User's Data	32
Preformatting Control Areas	32
Backing Up Data	32
Exporting and Importing a Data Set	33
Making a Copy of a Data Set	33
Ensuring Accessibility of Secondary Extents	33
Corrective Measures	34
Protecting VSAM Catalogs	34
Using REPRO to Back Up a Catalog	35
Backing up the Master Catalog	35
Dumping a Catalog and Its Data Sets	35
Updating a Backup Catalog	36
VSAM Volume Cleanup	37
Changing a Volume's Serial Number	38
VSAM Catalog Cleanup	38
Exporting a Catalog's Data Sets	39
Moving a Catalog to Another System	39
Catalog Recovery	40
Regaining Access to Data	40
Creating and Cataloging Objects	45
Preventing Duplicate Names in Catalogs	46
JCL and Dynamic Allocation (DEFINE)	46
Order of Catalog Use: DEFINE	47
The DEFINE Command: How Parameters are Organized	47
How to Use One Object as a Model for Another Object	49
Defining a Catalog	50
How Space Is Assigned to a Catalog	51
Identifying the Catalog's Volume	53
NonVSAM Data Sets and the Catalog's Update Password	54
Estimating the Catalog's Space Requirements	54
The Catalog's Secondary Allocation Amount	56

Space Allocation for the Catalog Recovery Area (CRA)	57
Defining a Data Space	57
VSAM Data Spaces on a Volume	58
Allocating the Volume	58
VSAM Objects in a Data Space	58
Space Assignment to VSAM Objects	59
Defining a Cluster	59
Specifying Cluster Information	60
Descriptive Information	60
Performance Options Information	61
Protection and Integrity Information	61
Defining an Alternate Index	62
Specifying Alternate-Index Information	62
Descriptive Information	63
Performance Options Information	63
Protection and Integrity Information	64
Defining a Path	64
Defining a NonVSAM Data Set	64
Defining an Alternate Name	65
How to Use an Alias to Identify a User Catalog	65
Defining a Generation Data Group	66
Defining a Page Space	66
Building an Alternate Index	68
How an Alternate Index is Built	68
DD Statements That Describe the SORT Workfiles	70
Order of Catalog Use: BLDINDEX	70
Modifying Catalog Information	72
Order of Catalog Use: ALTER	73
Generic Names and ALTER	73
Renaming Generically-Named Entries	74
VSAM Volume Recovery Function	74
Displaying Catalog Information	76
Order of Catalog Use: LISTCAT	76
Deleting Catalog Entries	78
Order of Catalog Use: DELETE	81
Generic Names and DELETE	82
Moving Entries	83
Exporting an Entry	84
Importing an Entry	85
Converting an OS Catalog's Entries to VSAM Catalog Entries	87
Connecting an OS CVOL to the Master Catalog	87
Using the DEFINE Command to Convert Some of the CVOL's	87
Entries	87
Using the CNVTCAT Command to Convert All of the OS CVOL's	88
Entries to VSAM Catalog Entries	88
Order of Catalog Use: CNVTCAT	91
Restoring a Cluster's End-of-File Values	92
Restoring Catalog Entries after System Failure	93
Listing the Catalog Recovery Area's Contents	94
Copying a Catalog Entry From the Catalog Recovery Area	96
Restoring the Catalog Entry That Was Obtained Using the	
EXPORTRA Command	98
Resetting Catalog Entries	99
RESETCAT Requirements	101
WORKFILE Space Requirements	101
Considerations for Multivolume Data Sets	102
JCL Requirements	103
Copying and Printing	105
Copying Data Sets	105
Copying a Catalog	108
Copy-Catalog Preparation	108
Copy-Catalog Procedure	110
Backing Up a Catalog	112
Unloading a Catalog	112
Reloading a Catalog	113

Optimizing the Performance of Unload/Reload	114
Printing Data Sets	115
Listing Tape Volumes Mounted at Checkpoint	117
Command Format	118
Functional Command Format	118
ALTER	120
ALTER Parameters: Summary	120
Entry-Types That Can Be Altered	121
ALTER Parameters	124
Required Parameters	124
Optional Parameters	124
ALTER Examples	137
Alter a Cluster's Entry: Example 1	137
Alter the Entrynames of Generically Named Clusters: Example 2	137
Alter the Attributes of a Generation Data Group: Example 3	138
BLDINDEX	139
BLDINDEX Parameters	139
Required Parameters	139
Optional Parameters	140
BLDINDEX Example	142
CHKLIST	143
CHKLIST Parameters	143
Required Parameter	143
Optional Parameters	143
CHKLIST Examples	144
Selecting Specific Checkpoint Entries: Example 1	144
Partitioned Checkpoint Data Set: Example 2	144
CNVTCAT	146
CNVTCAT Parameters	146
Required Parameters	146
Optional Parameters	146
CNVTCAT Examples	148
Convert an OS Catalog's Entries to Entries in a VSAM Catalog: Example 1	148
Convert an OS Catalog's Entries to Entries in a VSAM Catalog: Example 2	149
DEFINE ALIAS	150
DEFINE ALIAS Parameters	150
Required Parameters	150
Optional Parameter	150
DEFINE ALIAS Examples	152
Define an Alias for a NonVSAM Data Set: Example 1	152
Define an Alias for a User Catalog: Example 2	152
Define ALTERNATEINDEX	153
Define ALTERNATEINDEX	154
DEFINE ALTERNATEINDEX Parameters: Summary	156
DEFINE ALTERNATEINDEX Parameters	157
Required Parameters	157
Optional Parameters	161
Components (Alternate Index)	177
Catalog (Alternate Index)	180
DEFINE ALTERNATEINDEX Example	181
Define an Alternate Index: Example 1	181
DEFINE CLUSTER	183
DEFINE CLUSTER	184
DEFINE CLUSTER Parameters: Summary	186
DEFINE CLUSTER Parameters	187
Required Parameters	187
Data Organization	190
Optional Parameters	191
Components (Cluster)	206
Catalog (Cluster)	209
DEFINE CLUSTER Examples	211
Define a Key-Sequenced Cluster: Example 1	211
Define a Key-Sequenced Cluster and an Entry-Sequenced Cluster: Example 2	212
Define a Key-Sequenced Cluster (In a Unique Data Space): Example 3	213
Define a Relative-Record Cluster: Example 4	214
Define a Reusable Entry-Sequenced Cluster: Example 5	215

Define a Key-Sequenced Cluster: Example 6	216
Define an Entry-Sequenced Cluster (The Cluster has a Generic Name): Example 7	217
DEFINE GENERATIONDATAGROUP	219
DEFINE GENERATIONDATAGROUP Parameters	219
Required Parameters	219
Optional Parameters	219
Catalog (Generation Data Group)	221
DEFINE GENERATIONDATAGROUP Example	222
Define a Generation Data Group and a Generation Data Set Within It: Example 1	222
DEFINE NONVSAM	224
DEFINE NONVSAM Parameters	224
Required Parameters	224
Optional Parameters	226
Catalog (NonVSAM)	227
DEFINE NONVSAM Example	228
Define a NonVSAM Data Set: Example 1	228
DEFINE PAGESPACE	229
DEFINE PAGESPACE Parameters: Summary	229
DEFINE PAGESPACE Parameters	230
Required Parameters	230
Optional Parameters	231
Catalog (Page Space)	238
DEFINE PAGESPACE Examples	239
Define a Conventional (NOSWAP) Page Space: Example 1	239
Define a SWAP Page Space: Example 2	239
DEFINE PATH	241
DEFINE PATH Parameters: Summary	241
DEFINE PATH Parameters	242
Required Parameters	242
Optional Parameters	242
Catalog (Path)	248
DEFINE PATH Example	249
Define a Path: Example 1	249
DEFINE SPACE	250
DEFINE SPACE Parameters	250
Required Parameters	250
Optional Parameters	252
Catalog (Space)	252
DEFINE SPACE Example	253
Define a Data Space: Example 1	253
DEFINE USERCATALOG	254
DEFINE USERCATALOG Parameters: Summary	254
DEFINE USERCATALOG Parameters	255
Required Parameters	255
Optional Parameters	257
Components (Catalog)	263
Catalog (Catalog)	264
DEFINE USERCATALOG Examples	265
Defining a User Catalog and Determining the Catalog's Space Requirements: Example 1	265
Define a User Catalog Leaving Available Suballocatable Space: Example 2	267
Define a User Catalog Using the MODEL Parameter: Example 3	268
Define a Recoverable User Catalog: Example 4	270
DELETE	272
DELETE Parameters	272
Required Parameters	272
Optional Parameters	273
DELETE Examples	279
Deleting a Key-Sequenced VSAM Cluster: Example 1	279
Deleting an Entry-Sequenced VSAM Cluster and All Empty Data Spaces On a Volume: Example 2	280
Deleting Two Key-Sequenced Clusters: Example 3	281
Deleting a NonVSAM Data Set's Entry: Example 4	281
Deleting Nonempty VSAM Data Spaces on a Volume: Example 5	282
Deleting a User Catalog: Example 6	283
Deleting an Alias Entry: Example 7	283
Deleting Generically-Named Entries: Example 8	284
List a Generation Data Group's Entries, Then Delete the Group and Its Data Sets: Example 9	285

Deleting a Member of a Partitioned (NonVSAM) Data Set:	
Example 10	286
Deleting a Page Space: Example 11	287
EXPORT	288
EXPORT Parameters	288
Required Parameters	288
Optional Parameters	289
EXPORT Examples	291
Exporting a Key-Sequenced Cluster: Example 1	291
Exporting an Entry-Sequenced Cluster: Example 2	291
Exporting a User Catalog: Example 3	292
EXPORTRA	294
EXPORTRA Parameters	294
Required Parameters	294
Optional Parameters	296
EXPORTRA Examples	297
Using EXPORTRA for All Entries on One Volume: Example 1	297
Using EXPORTRA for All Entries on Multiple Volumes:	
Example 2	297
Using EXPORTRA for Selected Entries: Example 3	299
IMPORT	301
IMPORT Parameters	301
Required Parameters	301
Optional Parameters	303
IMPORT Examples	309
Importing a User Catalog: Example 1	309
Importing a Key-Sequenced Cluster: Example 2	309
Importing an Entry-Sequenced VSAM Cluster: Example 3	310
IMPORTRA	312
IMPORTRA Parameters	312
Required Parameters	312
Optional Parameters	312
IMPORTRA Example	316
Recovering a VSAM Catalog: Example 1—Part 1 (EXPORTRA)	316
Recovering a VSAM Catalog: Example 1—Part 2 (IMPORTRA)	317
LISTCAT	318
LISTCAT Parameters	318
Required Parameters	318
Optional Parameters	318
LISTCAT Examples	324
Listing a Key-Sequenced Cluster's Entry: Example 1	324
Alter a Catalog Entry, Then List the Modified Entry:	
Example 2	324
List Catalog Entries: Example 3	325
LISTCRA	326
LISTCRA Parameters	326
Required Parameters	326
Optional Parameters	326
LISTCRA Example	328
Listing a Catalog Recovery Area: Example 1	328
PRINT	329
PRINT Parameters	329
Required Parameters	329
Optional Parameters	329
PRINT Examples	334
Print a Key-Sequenced Cluster's Data Records: Example 1	334
Copy Records From a NonVSAM Data Set Into an	
Entry-Sequenced VSAM Cluster, Then Print the Records:	
Example 2	335
Print a VSAM Catalog: Example 3	337
REPRO	339
REPRO Parameters	339
Required Parameters	339
Optional Parameters	340
REPRO Examples	344
Copy Records Into a VSAM Cluster: Example 1	344
Copy Records Into a VSAM Cluster: Example 2	345
Copy a Catalog: Example 3	347
Unload a VSAM User Catalog: Example 4	349
Reload an Unloaded VSAM User Catalog: Example 5	350
RESETCAT	351
RESETCAT Parameters	351
Required Parameters	351

Optional Parameters	352
RESETCAT Examples	354
Resetting a Catalog: Example 1	354
Resetting a Catalog: Example 2	355
Resetting a Catalog: Example 3	356
VERIFY	358
VERIFY Parameters	358
Required Parameters	358
VERIFY Example	359
Upgrading a Data Set's End-of-File Information: Example 1	359
Controlling Command Execution	360
Condition Codes	360
IF-THEN-ELSE Command Sequence	361
DO-END Command Sequence	362
Null Commands	362
SET Command	362
PARM Command	363
Control Command Execution Examples	365
Control Command Execution: Example 1	365
Control Command Execution: Example 2	365
Control Command Execution: Example 3	365
Control Command Execution: Example 4	365
Control Command Execution: Example 5	366
Control Command Execution: Example 6	366
Appendix A. Examples of Jobs Using Access Method Services	
Commands	367
Example 1: Define a VSAM User Catalog	367
Explanation of Commands	367
Example 2: Define a VSAM User Catalog and a VSAM Data Space	368
Explanation of Commands	368
Example 3: Define VSAM Data Sets	369
Explanation of Commands	371
Example 4: Define NonVSAM and VSAM Data Sets	373
Explanation of Job Control Language Statements	375
Explanation of Commands	375
Example 5: Copying and Printing	377
Explanation of Job Control Language Statements	379
Explanation of Commands	379
Example 6: Record Replacement	381
Explanation of Job Control Language Statements	381
Explanation of Commands	381
Example 7: Alter the Cataloged Attributes of VSAM Data Sets	382
Explanation of Commands	383
Example 8: Creating an Alternate Index and Its Path	384
Explanation of Job Control Language Statements	384
Explanation of Commands	385
Example 9: Exporting VSAM Data Sets	386
Explanation of Job Control Language Statements	387
Explanation of Commands	387
Example 10: Importing VSAM Data Sets	389
Explanation of Job Control Language Statements	390
Explanation of Commands	391
Example 11: Deleting a VSAM Catalog and Its Cataloged Objects	392
Explanation of Commands	393
Appendix B. Interpreting LISTCAT Output Listings	396
LISTCAT Output Keywords	396
Alias Entry Keywords	397
Alternate-Index Entry Keywords	397
Cluster Entry Keywords	397
Data Entry Keywords	398
Index Entry Keywords	399
Generation Data Group Base Entry Keywords	401
NonVSAM Entry Keywords	401
Page Space Entry Keywords	401
Path Entry Keywords	402
User Catalog Entry Keywords	402
Volume Entry Keywords	402

Description of Keyword Fields	403
ALC: Allocation Group	403
ASN: Associations Group	404
ATT: Attributes Group	405
DSP: Data Space Group	408
GDG: Generation Data Group Base Entry, Special Fields For	409
NVS: NonVSAM Entry, Special Field For	409
HIS: History Group	410
PRT: Protection Group	410
STA: Statistics Group	411
VLS: Volumes Group	412
VOL: Volume Entry, Special Fields For	414
Device Type Translate Table	415
Examples of LISTCAT Output Listings	415
Job Control Language (JCL) for LISTCAT Jobs	415
LISTCAT and Access Method Services Output Messages	417
LISTCAT Output Listing	418
LISTCAT VOLUME Output Listing	420
LISTCAT SPACE ALL Output Listing	422
LISTCAT ALL Output Listing	423
LISTCAT ALLOCATION Output Listing	430
LISTCAT HISTORY Output Listing	433
LISTCAT CREATION/EXPIRATION Output Listing	435
Examples of LISTCAT in a TSO Environment	437
Appendix C. JCL DD Parameters to Take Care with	439
Appendix D. Interpreting LISTCRA Output Listings	442
Examples of LISTCRA Listings	443
Job Control Language (JCL) for LISTCRA Jobs	444
LISTCRA Output Listing	446
LISTCRA DUMP Output Listing	449
LISTCRA COMPARE Output Listing	452
LISTCRA DUMP COMPARE Output Listing	454
Appendix E. Sample Output from CHKLIST	455
Appendix F. Command Reference Summary	456
ALTER—modify attributes of a data set cataloged in a VSAM catalog	457
BLDINDEX—build alternate indexes for existing VSAM clusters	458
CHKLIST—identify tape volumes mounted when a checkpoint was taken	458
CNVTCAT—convert OS catalog entries into OS/VS2 VSAM catalog entries	458
DEFINE ALIAS—define an alternate name for a user catalog or a nonVSAM data set	458
Define Alternateindex—Define A VSAM Catalog Entry For An Alternate Index	459
Define Alternateindex (Continued)	460
DEFINE ALTERNATEINDEX (Continued)	461
DEFINE CLUSTER—define a VSAM catalog entry for a VSAM cluster	462
DEFINE CLUSTER (Continued)	463
DEFINE GENERATIONDATAGROUP—create a catalog entry for a generation data group	464
DEFINE NONVSAM—define a VSAM catalog entry for a nonVSAM data set	464
DEFINE PAGESPACE—define a VSAM catalog entry for a page space	464
DEFINE PATH—define a relationship between an alternate index and its base cluster	465
DEFINE SPACE—define a VSAM data space	465
DEFINE USERCATALOG—Create a VSAM user catalog	466
DELETE—delete entries from a VSAM catalog	467
EXPORT—disconnect a VSAM user catalog	467
EXPORT—copy a VSAM cluster or alternate index, together with its catalog entries, to a sequential data set	467
EXPORTRA—copy catalog entries from the Catalog Recovery Area (CRA)	468
IMPORT—connect a VSAM user catalog	468
IMPORT—restore a VSAM cluster or alternate index from a portable data set created by the EXPORT command	468

IMPORTRA—restore catalog entries from a portable data set created by the EXPORTRA command	469
LISTCAT—list entries from a catalog	469
LISTCRA—list the contents of the Catalog Recovery Area (CRA)	469
PRINT—print the contents of a data set	470
REPRO—copy data sets	470
RESETCAT—synchronize a recoverable catalog to the level of its owned volumes	470
VERIFY—restore a VSAM cluster's end-of-file values	470
Execution Control Commands	471
IF-THEN-ELSE Command Sequence	471
PARM Command	471
SET Command	471
Appendix G. Invoking Access Method Services from a Problem Program	472
Authorized Program Facility (APF)	472
Invoking Macro Instructions	472
LINK or ATTACH Macro Instruction	472
LOAD and CALL Macro Instructions	473
Invocation From a PL/I Program	474
Processor Invocation	475
Processor Condition Codes	477
User I/O Routines	477
Appendix H. Changing the Attributes of the VSAM Master Catalog	480
Glossary	486
Index	492

FIGURES

Figure 1.	Tasks and Commands	1
Figure 2.	Catalog Relationships in an OS/VS2 System	20
Figure 3.	Worksheet for Estimating a Catalog's Space Requirements	55
Figure 4.	OS Catalog Entry Types and VSAM Equivalents	89
Figure 5.	Converting Control Volume Entries	90
Figure 6.	Catalog Recovery Area Contents	93
Figure 7.	Adding Records To Various Types of Output Data Sets	107
Figure 8.	ALTER Parameters and the Entry-Types to Which Each Applies	122
Figure 9.	Completed Worksheet for Determining a Catalog's Space Requirements	266
Figure 10.	Character Format	330
Figure 11.	Dump Format	330
Figure 12.	Hex Format	330
Figure 13.	An Example of the Printed Record in DUMP Format	334
Figure 14.	An Example of the Printed Record in Hexadecimal	336
Figure 15.	An Example of a Printed Alphameric Character Record	336
Figure 16.	Messages that Follow the Entry Listing	417
Figure 17.	An Example of LISTCAT Output When No Parameters are Specified	419
Figure 18.	An Example of LISTCAT VOLUME Output	420
Figure 19.	An Example of LISTCAT SPACE ALL Output	422
Figure 20.	An Example of LISTCAT ALL Output	423
Figure 21.	An Example of LISTCAT ALLOCATION Output	430
Figure 22.	An Example of LISTCAT HISTORY Output	433
Figure 23.	An Example of LISTCAT CREATION/EXPIRATION Output	435
Figure 24.	JCL DD Parameters	439
Figure 25.	An Example of LISTCRA NAME NOCOMPARE Output	447
Figure 26.	An Example of LISTCRA DUMP NOCOMPARE Output	450
Figure 27.	An Example of LISTCRA NAME COMPARE Output	453
Figure 28.	An Example of LISTCRA DUMP COMPARE Output	454
Figure 29.	An Example of CHKLIST Output	455
Figure 30.	Processor Invocation Argument List From a Problem Program	476
Figure 31.	Arguments Passed to and From a User I/O Routine	479
Figure 32.	Sample Job Stream to Make the Master Catalog Recoverable	482

GUIDE TO ACCESS METHOD SERVICES

Figure 1 shows a list of tasks that Access Method Services commands can be used to perform. The left-hand column shows tasks that you might want to perform. The middle column more specifically defines the tasks. The right-hand column shows the commands that can be used to perform each task.

Operation		Command
Alter	the information in a catalog entry	ALTER
Attach	a user catalog to the master catalog	DEFINE USERCATALOG or IMPORT CONNECT
Backup	a VSAM data set a nonVSAM data set a VSAM catalog	EXPORT or REPRO REPRO REPRO
Build	an alternate index a generation data group	DEFINE ALTERNATEINDEX and BLDINDEX DEFINE GENERATIONDATAGROUP
Catalog	a VSAM data set a nonVSAM data set	DEFINE CLUSTER and DEFINE ALTERNATE INDEX DEFINE NONVSAM
Change	a data set's description in the catalog the device type of the volume on which the catalog resides	ALTER REPRO
Connect	a user catalog to a master catalog an OS/VS CVOL to a master catalog	IMPORT CONNECT DEFINE NONVSAM
Convert	a data set to VSAM format a VSAM data set to sequential format OS/VS CVOL entries to VSAM catalog entries	REPRO REPRO CNVTCAT
Copy	a data set or catalog	REPRO
Create	a backup copy of a data set a catalog an alias for a nonVSAM data set an alias for a VSAM data set a VSAM data set an alternate index the relationship between an alternate index and its base cluster a generation data group a generation data set a page space a VSAM data space a catalog entry for a nonVSAM data set	REPRO or EXPORT DEFINE USERCATALOG DEFINE ALIAS DEFINE PATH DEFINE CLUSTER DEFINE ALTERNATEINDEX and BLDINDEX DEFINE PATH DEFINE GENERATIONDATAGROUP DEFINE NONVSAM DEFINE PAGESPACE DEFINE SPACE DEFINE NONVSAM
Define	(see above entries for Create)	

Figure 1 (Part 1 of 3). Tasks and Commands

Operation		Command	
Delete	a user catalog	DELETE USERCATALOG	
	an alias	DELETE ALIAS	
	an alternate index and its paths	DELETE ALTERNATEINDEX	
	a path	DELETE PATH	
	any type of catalog entry	DELETE	
	a VSAM data set	DELETE CLUSTER	
	a VSAM data space	DELETE SPACE	
	a generation data group	DELETE	
	a nonVSAM data set	GENERATIONDATAGROUP	
	a page space	DELETE NONVSAM DELETE PAGESPACE	
Disconnect	a user catalog	EXPORT DISCONNECT	
Display	a catalog's contents	LISTCAT	
	a catalog recovery area's contents	LISTCRA	
Enter	a nonVSAM data set in a catalog	DEFINE NONVSAM	
Export	a user catalog	EXPORT DISCONNECT	
	a VSAM data set	EXPORT or EXPORTRA	
Import	a user catalog	IMPORT CONNECT	
	a VSAM data set	IMPORT or IMPORTRA	
List	a password	LISTCAT	
	a data set's contents	PRINT	
	contents of a catalog or of a catalog entry	LISTCAT	
	contents of a catalog recovery area	LISTCRA	
	tapes mounted at a checkpoint	CHKLIST	
Load	records into a data set	REPRO	
	an alternate index	BLDINDEX	
Modify	a data set's description in the catalog	ALTER	
Move	a catalog to another system	EXPORT DISCONNECT and IMPORT CONNECT	
	a VSAM data set to another system	EXPORT and IMPORT	
	a nonVSAM data set to another system	REPRO	
Password Protect	establish for a VSAM data set	DEFINE CLUSTER	
	alternate index	DEFINE ALTERNATEINDEX	
	path	DEFINE PATH	
	catalog	DEFINE USERCATALOG or DEFINE MASTERCATALOG	
	page space	DEFINE PAGESPACE	
	modify an existing password, or add a password to an existing catalog entry	ALTER	
	delete a password	ALTER	
	list passwords	LISTCAT	
	Print	a data set	PRINT
	Recover	from a processing-program failure	VERIFY
from a catalog failure		LISTCRA and EXPORTRA and IMPORTRA	
		RESETCAT	

Figure 1 (Part 2 of 3). Tasks and Commands

Operation		Command
Release	a user catalog from the master catalog	EXPORT DISCONNECT
Rename	a data set	ALTER
Reproduce	a VSAM or nonVSAM data set	REPRO
Restore	a data set's end-of-file information a catalog entry and/or the contents of its object	VERIFY EXPORTRA and IMPORTRA RESETCAT
Uncatalog	a data set	DELETE
Unload	a data set or catalog	REPRO
Verify	a VSAM data set's end-of-file indicators	VERIFY

Figure 1 (Part 3 of 3). Tasks and Commands

INTRODUCTION

Access Method Services is a service program that is used with VSAM (Virtual Storage Access Method) to establish and maintain catalogs and data sets. If you plan on using VSAM, or if you are responsible for maintaining the system catalog, you must use Access Method Services commands.

This chapter introduces the commands available through Access Method Services, provides the background information about VSAM that is required to establish and maintain data sets and catalogs, and briefly describes the use of the Time Sharing Option (TSO) with VSAM and Access Method Services.

ACCESS METHOD SERVICES COMMANDS

Access Method Services enables you to define a VSAM data set and load records into it, convert a sequential or an indexed-sequential data set to the VSAM format, list VSAM catalog information or data-set records, copy a data set for reorganization, create a backup copy of a data set or catalog, convert an OS CVOL's entries to VSAM catalog entries, and make a data set portable from one operating system to another.

The user invokes Access Method Services functions by issuing a command and specifying its parameters. The user can execute the IDCAMS program and include the command and its parameters as input to the program. The user can also call the IDCAMS program from within another program and pass the command and its parameters to the IDCAMS program. TSO (Time Sharing Option) users can execute Access Method Services functional commands from the TSO terminal as though they were TSO commands.

There are Access Method Services commands for:

- Defining, altering, and deleting data sets
- Listing catalog entries
- Copying and printing data sets
- Constructing alternate indexes for VSAM data sets
- Moving catalogs and data sets from one operating system to another
- Aiding in recovery from damage to data sets or catalogs
- Converting an OS catalog's entries to VSAM catalog entries
- Listing tape volumes that were mounted at the time of a checkpoint
- Controlling command execution by testing or setting condition codes
- Establishing diagnostic-aids and printed-output options

There are two types of Access Method Services commands: functional commands that are used to request the actual work—for example, defining a data set or listing a catalog—and modal commands that allow the conditional execution of functional commands. TSO users are allowed to use only the functional commands.

FUNCTIONAL COMMANDS

The functional commands are:

- ALTER, which is used to alter previously defined catalog entries.
- BLDINDEX, which constructs alternate indexes for existing data sets.
- CHKLIST, which is used to identify tape volumes mounted when a checkpoint was taken.
- CNVTCAT, which is used to convert entries in an OS CVOL (control volume) into entries in a VSAM catalog.
- DEFINE, which is used to create catalog entries for data sets, catalogs, and space that VSAM is to allocate from.
- DELETE, which is used to delete catalog entries.
- EXPORT, which is used to create a copy of a VSAM data set for backup or to make a data set or user catalog portable so that it can be used on another system.
- EXPORTRA, which retrieves VSAM data sets and catalog entries that are no longer accessible from a VSAM catalog.
- IMPORT, which is used to read a backup copy of a VSAM data set or to make a data set or catalog that was previously exported from one system available for use in another system.
- IMPORTRA, which returns VSAM data sets and catalog entries to an accessible condition.
- LISTCAT, which is used to list catalog entries.
- LISTCRA, which is used to diagnose suspected problems in VSAM catalogs.
- PRINT, which is used to print VSAM, ISAM, or SAM data sets.
- REPRO, which is used to copy data sets, to convert sequential and indexed-sequential data sets to VSAM format, to convert VSAM and indexed-sequential data sets to sequential format, to back up a VSAM catalog, to read a backup copy of a VSAM catalog, and to copy VSAM catalogs.
- RESETCAT, which synchronizes a catalog to the level of its owned volumes.
- VERIFY, which is used to cause a catalog to correctly reflect the end of a data set after an error occurred in closing a VSAM data set that may have caused the catalog to be incorrect.

The functional commands that can be used on nonVSAM data sets include DEFINE, ALTER, DELETE, LISTCAT, CNVTCAT, REPRO, and PRINT.

The functional commands that can be used on page spaces include DEFINE, DELETE, ALTER, and LISTCAT.

All functional commands except CHKLIST can be used when you execute Access Method Services at a TSO terminal.

MODAL COMMANDS

The modal commands, ones that control command execution and establish options, are:

- IF, which tests a condition code and executes according to the results of the test. IF is followed by THEN and ELSE clauses which specify alternative actions.
- DO-END, which denote the beginning and ending of a command sequence.
- SET, which changes condition codes.
- PARM, which specifies diagnostic-aids and printed-output options.

LANGUAGE CONSIDERATIONS

All Access Method Services commands have this general structure:

COMMAND parameters ... terminator

COMMAND specifies the type of service requested. The parameters further describe the service requested. Terminator indicates the end of the command statement.

Commands can begin at or to the right of the left margin. The default margins are 2 and 72 for batch processing jobs. Commands are separated from their parameters by one or more separators, that is, one or more blanks, commas, or comments. Comments are strings of characters surrounded by a /* and an */. Comments can contain any characters you desire except an '*/'.

Many of the commands can be abbreviated. The abbreviations that are allowed are listed in the discussion of each command as well as in Appendix F. Restrictions that apply to abbreviations under TSO are described in "Time Sharing Option (TSO)."

Parameter Set

A parameter can be either a positional parameter or a keyword parameter. A positional parameter is characterized by its position in relation to other parameters. In Access Method Services, positional parameters are always required. A keyword parameter is a value preceded by a specific character string. For example, in:

VOLUME (25DATA)

VOLUME is a keyword that indicates that the value 25DATA is a volume serial number.

A keyword parameter might have a set of subparameters. The subparameters must follow the same rules as parameter sets in general. When the subparameters are positional, the first subparameter is always required and the second may be optional (that is, you cannot specify the second subparameter without also specifying the first).

Many of the keywords may be abbreviated. The abbreviations permitted are listed with each parameter as well as in Appendix F. Some keywords are plural in form; they can also be coded in the singular form.

As a group, positional parameters must always appear first in a parameter set. Keyword parameters always follow any positional parameters. The order of keyword parameters is not important.

A parameter or subparameter can consist of a list of similar items. Positional parameter lists must be enclosed in parentheses unless the list contains only one item. The

parentheses can be preceded and followed by blanks, commas, or comments. For example:

```
DELETE(entryname [...])
```

indicates that if more than one entry is to be deleted, the list of entrynames must be enclosed in parentheses. However, if only one entryname is specified, the parentheses are not required.

An item in a list can be a parameter set itself. Each such item, as well as the list of items, is enclosed in parentheses. Given:

```
OBJECTS((entryname NEWNAME (newname))...)
```

the following are valid:

```
OBJECTS( -  
  ENTRY1 NEWNAME(NEWNAME1))
```

```
OBJECTS( -  
  (ENTRY1 NEWNAME(NEWNAME1)) -  
  (ENTRY2 NEWNAME(NEWNAME2)))
```

In the first case, only one entry is to be renamed. The entryname and its new name are enclosed in parentheses. In the second case, each entryname and its new name are enclosed in parentheses and the entire list is enclosed in parentheses.

All parameters and subparameters must be separated from one another by one or more separators (commas, blanks, or comments). The only exception is that parameters that immediately follow a subparameter set that is enclosed in parentheses do not need to be separated from the closing parenthesis.

The values you specify in the parameters can be surrounded by separators. Some values can be longer than a single record. When a value is longer than a single record, you indicate that it is continued by coding a plus sign (+) followed only by blanks or a comment. The first nonseparator character found in a record following the plus sign is treated as a continuation of the value.

A value cannot contain commas, semicolons, blanks, parentheses, or slashes unless the entire value is enclosed in single quotation marks. A single quotation mark in a field enclosed in single quotation marks must be coded as two single quotation marks.

In some parameters it is necessary to specify a password following the name of a catalog entry or the name of a JCL DD statement. You do this by coding the name, a slash, and the password. The slash can be surrounded by separators. For example:

```
DELETE PAYROLL/CTLGPAY
```

specifies the password CTLGPAY for the data set PAYROLL.

When you specify a password following a name, you must remember the convention for entering commands. If you code:

```
MYNAME/*WORD*/
```

MYNAME is treated as a name followed by the comment WORD. If you want *WORD*/ to be the password, you code either:

```
MYNAME/'*WORD*/'
```

or

```
MYNAME/ *WORD*/
```

Notice that the second example contains a blank after the first slash that separates the name from the password.

Continuing Commands

Commands can be continued on several records or lines. Each record or line except the last must have a hyphen or a plus sign as the last nonblank character before or at the right margin. A hyphen indicates continuation of the command. A plus sign indicates continuation of the command and continuation of a value within the command. These two types of continuation are shown in the example below:

```
DELETE -  
  (ENTRY1 -  
   ENTRY2 -  
   ENTR+  
     Y3) -  
NONVSAM
```

Blank records or records ending with complete comments must end with a continuation mark when they appear in the middle of a command and when they appear preceding or following the THEN and ELSE clauses of an IF command. Records ending with partial comments must always end with a continuation mark. Also, remember that only blank characters may appear between a continuation mark and the end of the record.

Note: The DO-END sequence does not require continuation characters.

Continuation Cautions

The continuation rules must be used cautiously when modal commands, comments, or blank records appear in the input stream. Blank records or records ending with complete comments must end with a continuation mark when these types of records appear in the middle of a command or when they appear preceding or following the THEN and ELSE clauses of an IF command. Records ending with partial comments must always end with a continuation mark.

You must be careful when continuing modal commands so that you don't inadvertently specify a null command. For information on null commands, see "Null Commands."

Remember that only blanks can appear between a continuation mark and the end of the record.

Note: The DO-END sequence does not require continuation characters; if they are used, they may be taken as a null command or cause unpredictable results.

The following examples show common continuation errors:

- IF LASTCC = 0 -
 THEN
 LISTCAT

A continuation mark (hyphen) is missing after the THEN keyword. A null command is assumed after the THEN keyword, and the LISTCAT command is unconditionally executed.

- IF LASTCC = 0 -
 THEN -
 REPRO ...
 /*ALTERNATE PATH*/
 ELSE -
 PRINT ...

Because no continuation mark (hyphen) follows the comment, a null command is assumed. The ELSE keyword will not match up with the THEN keyword. Note the correct use of the continuation marks on the other records.

- IF LASTCC = 0 -


```
THEN -  
REPRO ...  
ELSE -
```

```
PRINT ...
```

Because a blank line with no continuation mark (hyphen) follows the ELSE keyword, the ELSE becomes null and the PRINT command is unconditionally executed.

- PARM TEST (- /*COMMENT*/
TRACE)

The PARM command will not be continued onto the second record because characters other than blanks appear between the continuation mark (hyphen) and the end of the record.

- PARM TEST (TRA+
/*FIELD CONTINUATION*/
CE)

The end of the PARM command is found after the second record because no continuation was indicated. The command is rejected.

Terminator

The terminator indicates the end of the command. The terminator can be either an end of command condition (that is, no continuation mark) or a semicolon. If you use the semicolon as the terminator, the semicolon cannot be enclosed in quotation marks or embedded in a comment. Everything to the right of the semicolon is ignored. If there is information to the right of the semicolon that is continued to another record, all of the information including the continued information is ignored.

For example, if you coded:

```
PARM TEST (TRACE); PARM -  
GRAPHICS (CHAIN(TN))/*COMMENT*/  
PRINT ...  
REPRO ...
```

Characters following the semicolon terminator are ignored. The continuation mark (hyphen) at the end of the second record causes the PRINT command to also be ignored. The first PARM command and the REPRO command are the only commands that are recognized.

JCL AND DYNAMIC ALLOCATION

Access Method Services commands deal with both data sets and volumes which must be identified when used. The following is a summary of the data set and volume identification requirements:

- A data set must be identified when it is actually accessed. The following commands require that a specific data set or data sets be identified.
 - ALTER (when altering a member name in a partitioned data set)
 - BLDINDEX
 - DELETE (when the ERASE option is invoked, when deleting a member of a partitioned data set)
 - EXPORT
 - IMPORT

- PRINT
- REPRO
- VERIFY
- A volume or volumes must be identified when it is necessary for VSAM to access the Volume Table of Contents (VTOC), allocate or release space using OS/VS DADSM functions, or access a catalog recovery area (see the section "Catalog Recovery" for information on this option). The following commands require volume identification:
 - ALTER (for recoverable catalogs, for VTOC access in renaming a nonVSAM or VSAM unique data set)
 - DEFINE (for recoverable catalogs, for space allocation in defining catalogs, VSAM space, and VSAM unique data sets)
 - DELETE (for recoverable catalogs, for space deallocation in deletion of catalogs, VSAM space, and VSAM unique data sets, for deletion of nonVSAM data sets)

Other commands such as EXPORTRA, IMPORTRA, LISTCRA, and RESETCAT require volume identification for specific purposes. These requirements are covered in the discussion of each command.

A VSAM data set or volume can be identified through JCL or by the data-set name or volume serial number within the command that requires the data set or volume for its execution. If JCL is not used, the data set or volume is dynamically allocated, as required.

In order to dynamically allocate a VSAM data set, the data-set name must exist and be cataloged. The catalog that contains the entry must be either a user catalog identified with a JOBCAT or STEPCAT DD statement, a user catalog whose name or alias is the first qualifier of the qualified data-set name, or the master catalog.

In order to dynamically allocate a nonVSAM dataset, the data set name must exist and must be cataloged. The catalog that contains the entry must be either a VSAM user catalog identified with a JOBCAT or STEPCAT DD statement, a VSAM user catalog or OS CVOL whose name or alias is the first qualifier of the qualified data-set name, or the master catalog.

Access Method Services dynamically allocates VSAM and nonVSAM data sets with a disposition of OLD.

When a user catalog is identified by name in an Access Method Services command (using the CATALOG parameter) or when VSAM determines that a specific user catalog must be accessed, the user catalog is dynamically allocated if it has not been identified with a JOBCAT or STEPCAT DD statement.

In order to dynamically allocate a volume, it must already be mounted as permanently resident or reserved. The PRIVATE and PUBLIC use attributes should be carefully considered when you mount a volume.

When a JCL DD statement is used to identify a data set, the following information must be included on the DD statement.

- The data-set name
- The unit and volume serial number(s), if the data set is not cataloged
- The disposition

- For VSAM data sets, the AMP='AMORG' parameter is required if the JCL also supplies unit and volume serial number information

For example, the two DD statements shown below demonstrate two different methods by which you can describe and allocate a VSAM data set:

```
//VSAMDS1 DD DSN=VSAM.DATA.SET1,DISP=OLD
//VSAMDS2 DD DSN=VSAM.DATA.SET2,VOL=SER=VSER01,
// UNIT=3330,DISP=OLD,AMP='AMORG'
```

Access Method Services does not provide specific protection for data sets in a shared environment. Therefore, you should specify DISP=OLD on the DD statement for any data set which could be accessed improperly in a shared environment.

Access Method Services uses default buffering for VSAM data sets; that is, the buffer space value contained in the data set's catalog entry. You may wish to override this value using the BUFND, BUFNI, or BUFSP subparameter of the AMP parameter (see OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for information on these parameters).

To identify and allocate a volume, the following information must be included:

- The volume serial number
- The unit

For example, the DD statement shown below identifies and allocates a VSAM volume:

```
//VOLDD DD VOL=SER=VSER01,UNIT=3330,DISP=OLD
```

Concatenated DD statements are supported under circumstances explicitly specified in the remainder of this publication.

The section "JOB CAT and STEP CAT DD Statements" contains a description of the DD statements required for user catalogs.

Certain Access Method Services commands use nonVSAM data sets. They are: EXPORT, IMPORT, EXPORTRA, and IMPORTRA. The examples provided for each command in the section "Command Format," as well as the examples in Appendix A, show the DD statements for the required nonVSAM data sets.

If the Access Method Services processor encounters a condition which requires it to abort a job, it takes a snap dump of virtual storage. An AMSDUMP DD statement is required to obtain the snap dump; for example,

```
//AMSDUMP DD SYSOUT=A
```

If you do not supply an AMSDUMP DD statement and the processor encounters a condition requiring the job to be aborted, it produces an abbreviated dump only.

OUTPUT DATA SETS

The normal output data set for listing is SYS PRINT. The default parameters of this data set are:

- Record format: VBA
- Logical record length: 125, that is, 121+4
- Blocksize: 629, that is, 5x(121+4)+4

Print lines are 121 bytes in length. The first byte is the ANSI control character. The minimum specifiable LRECL is 121

(U-format records only). If a smaller size is specified, it is overridden to 121.

It is possible to alter the above defaults through specification of the desired values in the DCB parameter of the SYSPRINT statement. The record format, however, cannot be specified as F or FB. If you do specify either one, it is changed to VBA.

In several commands you have the option of specifying an alternate output data set for listing. If you do specify an alternate, you must specify DCB parameters in the referenced DD statement using the attributes described above. When specifying an alternate output data set, you should not specify F or FB record formats. Note, however, that JCL statements, system messages, and job statistics are written to the SYSPRINT output device, not to the alternate output data set. "Appendix D Interpreting LISTCRA Output Listings" shows an example of the use of an alternate output data set.

INVOKING ACCESS METHOD SERVICES

When you want to use an Access Method Services function, you invoke the Access Method Services processor. The processor decodes your request (that is, the list of commands and parameters you supply) one command at a time, then calls the appropriate functional routines to perform all services required by that command.

There are three ways you can invoke the Access Method Services processor:

- As a job or jobstep
- From a TSO terminal
- From within your own program. (See "Appendix G: Invoking Access Method Services from a Problem Program.")

If you want more details on the structure and operation of Access Method Services in an OS/VS2 system, see OS/VS2 Access Method Services Logic.

AS A JOB OR JOBSTEP

You can use job control language (JCL) statements to invoke the Access Method Services processor. The examples that illustrate each command's use invoke the processor with JCL statements. You identify the Access Method Services processor with PGM=IDCAMS.

```
//YOURJOB JOB YOUR INSTALLATION'S JOB=ACCOUNTING DATA
//JOB CAT DD DSNAME=YOUR.CATALOG,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//STEP CAT DD DSNAME=ANOTHER.CATALOG,DISP=SHR
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
```

Access Method Services commands and their parameters

/*

- //YOURJOB is required. The JOB statement describes your job to the OS/VS2 system. Your installation might require you to supply user identification, accounting, and authorization information with the JOB statement's parameters.
- //JOB CAT is optional. The JOB CAT DD statement identifies a user catalog that can be used by each of the job's steps. Because the master catalog is always open and available to all jobs on the system, you should not identify it with a JOB CAT DD statement.

- //STEP1 is required (that is, an EXEC statement is required). The EXEC statement invokes the Access Method Services processor to decode and process the input statements (that is, the Access Method Services commands and parameters contained in the SYSIN data set). You can use the PARM operand of the EXEC statement to pass parameters to the Access Method Services processor. See the description of the PARM command in the chapter "Controlling Command Execution" for an explanation of the options you can specify.
- //STEPCAT is optional. The STEPCAT DD statement identifies a user catalog that can be used when processing the jobstep. If user catalogs are identified with JOBCAT and STEPCAT DD statements, only the catalog(s) identified with the STEPCAT DD statement and the master catalog are used with the jobstep. Because the master catalog is always open and available to all jobs on the system, you should not identify it with a STEPCAT DD statement.
- //SYSPRINT is required. The SYSPRINT DD statement identifies the output device (usually a printer in a batch job) that Access Method Services sends messages and output information to. (See "Output Data Sets" for more details on how to describe an output device other than SYSOUT=A.)
- //SYSIN is required. The SYSIN DD statement identifies the source of the input statements. An input statement, to Access Method Services, is a functional or modal command and its parameters. When you code SYSIN DD *, you identify the statements immediately following as input.

The last input statement may be followed by a delimiter statement, which has '/'* in the first two columns. The restrictions that apply to all Access Method Services commands are described in the section "Language Considerations," above.

FROM A TSO TERMINAL

When you use the Time Sharing Option (TSO) to process your data, you can invoke Access Method Services from your TSO terminal. Each time you issue an Access Method Services command as a TSO command, TSO builds the appropriate interface information and invokes Access Method Services.

You can issue one command at a time. Access Method Services processes the command completely before TSO allows you to continue processing. The following Access Method Services commands are supported in a TSO environment:

```

ALTER
BLDINDEX
CNVTCAT
DEFINE -
  ALIAS
  ALTERNATEINDEX
  CLUSTER
  GENERATIONDATAGROUP
  NONVSAM
  PAGESPACE
  PATH
  SPACE
  USERCATALOG.
DELETE
EXPORT
EXPORTRA
IMPORT
IMPORTRA
LISTCAT
LISTCRA
PRINT
REPRO
RESETCAT
VERIFY

```

The command you issue is processed immediately. You cannot issue modal commands (that is, the IF, THEN, ELSE, DO, END, PARM, and SET commands) that modify or make conditional the processing of subsequent functional commands.

The restrictions that apply to the coding of Access Method Services commands are described in the "Language Considerations" and "Time Sharing Option (TSO)" sections. Other TSO restrictions are noted with the descriptions of each appropriate parameter.

FROM A USER'S PROGRAM

A processing program can invoke Access Method Services with the ATTACH, LINK or LOAD and CALL macros. Before the program issues the invoking macro, however, it must initialize the appropriate register and argument list contents.

The register contents follow standard linkage conventions; that is, register 1 contains the address of the argument list, register 13 contains the address of a save area, register 14 contains the address of the return point, and register 15 contains the address of the entry point IDCAMS in Access Method Services.

The contents of the argument list are described in Appendix G.

AUTHORIZED PROGRAM FACILITY

The authorized program facility (APF) limits the use of sensitive system services and resources to authorized system and user programs. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor, for information about program authorization.

All Access Method Services load modules are contained in SYS1.LINKLIB, and the root segment load module (IDCAMS) is link edited with the SETCODE AC(1) attribute. These two characteristics ensure that Access Method Services executes with APF authorization.

APF authorization is established at the job step task level. If, during the execution of an APF-authorized job step, a load request is satisfied from an unauthorized library, the task will be abnormally terminated. It is the installation's responsibility to ensure that a load request cannot be satisfied

from an unauthorized library during Access Method Services processing.

The following situations could cause the APF authorization of Access Method Services to be violated:

- An Access Method Services module is loaded from an unauthorized library.
- A user security verification routine (USVR) is loaded from an unauthorized library during Access Method Services processing.
- An Exception Exit routine is loaded from an unauthorized library during Access Method Services processing.
- A user-supplied special graphics table is loaded from an unauthorized library during Access Method Services processing.

Since APF authorization is established at the job step task level, Access Method Services will not be authorized if invoked by an unauthorized problem program or an unauthorized Terminal Monitor Program (TMP).

Under TSO, if your system has neither OS/VS2 TSO/VTAM Selectable Unit (VS2.03.813) or the OS/VS2 MVS TSO Command Package Program Product, you can authorize your TMP by relink-editing it with the SETCODE AC(1) attribute. With either the OS/VS2 TSO/VTAM Selectable Unit or the OS/VS2 MVS TSO Command Package Program Product you must enter the names of those Access Method Services commands requiring APF authorization which you wish to execute under TSO in the authorized command list. If you have the OS/VS2 TSO/VTAM Selectable Unit, see OS/VS2 System Programming Library: TSO for information on how to enter the command names in the list. If you have the OS/VS2 MVS TSO Command Package Program Product, see OS/VS2 MVS TSO Command Package: User's Guide and Reference for the same information.

The restricted functions performed by Access Method Services that cannot be requested in an unauthorized state are:

- REPRO—the copy catalog facility, or the catalog unload/reload facility.
- PRINT—when the object to be printed is a VSAM catalog.
- LISTCRA
- EXPORTRA
- RESETCAT

If the above functions are required and Access Method Services is invoked from a problem program or a TSO Terminal Monitor Program, the invoking program must be authorized.

VSAM DATA SETS

You need to know about the structure and treatment of VSAM data sets to use Access Method Services. VSAM data can be stored in key sequence, entry sequence, or record-number sequence.

Records in a key-sequenced data set are stored in the order defined by the collating sequence of the contents of the key field in each record. Each record has a unique value, such as employee number or invoice number, in the key field. To determine where to insert a new record, VSAM uses an index that pairs the key of a record with the record's location.

Records in an entry-sequenced data set are stored without respect to the contents of the records. Their sequence is determined by the order in which they are stored: their entry sequence. A new record is stored after the last record in the data set.

Records in a relative-record data set are stored without respect to their contents and to their sequence of entry. Their sequence is determined by the user-specified record number that identifies the record's position in the data set. Each record is fixed-length and is stored in a slot. Records can be inserted between other records (that is, in an empty slot between two filled slots), can be added to the end of the data set, or can replace records within the data set.

When a data set is created, it is defined, along with its index, if any, in a cluster. A key-sequenced data set and its index make up a cluster. An entry-sequenced or relative record data set is also defined as a cluster, even though it does not have an index.

Apart from the primary index, one or more alternate indexes can be built over a single key-sequenced or entry-sequenced data set. Each alternate index accesses the data records of a given data set via a different key field (the alternate key) within these records. A VSAM cluster that contains data records pointed to by an alternate index is called a base cluster.

In order to gain access to a base cluster via an alternate index, you must define a path between the alternate index and the base cluster. When a path is defined, you must specify the name of the alternate index which is to be considered as the entry for the path. The termination of the path is the base cluster to which the alternate index is related. The path, which is an entry in the catalog, has a name of its own which always refers to the alternate index and the related base cluster as a pair.

In addition to data sets, VSAM accesses catalogs. All data—ordinary user data, indexes, and catalogs—is physically stored and manipulated in the same way.

Refer to Planning for Enhanced VSAM Under OS/VS for more information on the structure of VSAM data sets and how data is physically stored.

VSAM VOLUME OWNERSHIP

VSAM gets and manages space on a storage volume in units called data spaces. One or more data sets are stored in a data space. Conversely, a data set may be stored in one or more data spaces, on one or more storage volumes. Data sets and data spaces may be extended beyond their original size. Data spaces are extended by whole numbers of tracks or cylinders; VSAM extends them automatically as more space is needed. A data space can be as large as a volume and can have a maximum of 16 extents.

VSAM data spaces are defined in a VSAM catalog. That catalog controls (owns) the volumes that contain the data spaces defined in the catalog, including the volume that contains the catalog itself. The catalog also owns candidate volumes for VSAM data sets defined in the catalog. VSAM ownership of a volume is established with the DEFINE USERCATALOG or DEFINE SPACE commands or when the UNIQUE attribute is specified with DEFINE CLUSTER, DEFINE ALTERNATEINDEX, or DEFINE PAGESPACE.

All VSAM clusters, alternate indexes, page spaces, and data spaces stored on a volume must be cataloged in the catalog that owns the volume. On the other hand, nonVSAM data sets can be cataloged in any VSAM catalog.

VSAM's ownership and usage of space on a volume is indicated in the volume's table of contents (that is, in DSCBs in the volume's VTOC). Two data-set security bits in the format-1 (Identifier) DSCB of each VSAM data space on the volume are set to indicate that a password is required to read or write data in the data space. (They are set whether the object or objects in the data space are actually password-protected or not.) The ownership bit in the volume's format-4 (VTOC) DSCB is set to 1. The ownership bit indicates that the volume is owned by a VSAM catalog, but does not identify the owning catalog. Each VSAM catalog contains a volume entry for each volume it owns. The volume entry describes the volume's characteristics, each extent of the volume's VSAM data space(s), and each VSAM object that uses the volume's space.

In order to take away the ownership of a volume from a VSAM catalog, you must first delete all VSAM objects and then all data spaces on the volume. The DELETE SPACE command deletes the VSAM data spaces on the volume, removes the volume entry from the catalog, and revises the format-4 DSCB in the volume's VTOC. When you are unable to use the DELETE command because Access Method Services can no longer access the volume (due to damage that resulted from a system or hardware failure), you can reset the ownership bit by using the ALTER REMOVEVOLUMES command (see "Modifying Catalog Information").

The VTOC contains the name of each VSAM data space on the volume, and might contain VSAM-generated names for the data and index components of a cluster, alternate index, or page space. When you name the data or index component of a VSAM data set, alternate index, or page space, and when the object or its component is in its own data space (that is, it was defined with the UNIQUE attribute), the name of the component identifies the data space. Otherwise, VSAM generates a name for the data space. The name generated by VSAM has the following format:

- For a data space containing suballocated VSAM objects, the VSAM generated name is:

Z999999n.VSAMDSPC.Taaaaaaaa.Tbbbbbbb

where:

- n=2 if no catalog resides in the data space
 - n=4 if a user catalog resides in the data space, or if the master catalog resides in the data space and the master catalog was created under OS/VS2 Release 2 or a later release
 - n=6 if the master catalog resides in the data space
 - aaaaaaabbabbbb is the timestamp value
- For a unique data space (that is, a data space that cannot contain more than one cataloged VSAM object), the VSAM-generated name is:

VSAMDSET.Tbbbbbbb.DFDyyddd.Taaaaaaaa.Tbbbbbbb

where:

- yyddd is the date (year and Julian day)
- aaaaaaabbabbbb is the timestamp value

When you define a VSAM data set with the UNIQUE attribute, VSAM creates a unique data space. If you specify a name for the data and/or index component, VSAM places the name you specify in the format-1 DSCB rather than generating a name.

When you delete a VSAM data set with the UNIQUE attribute, the data space is also deleted (the volume must be mounted). The volume ownership bit is turned off; the volume entry in the

catalog is turned off.

If you do not specify a name for the data and index components of a suballocated VSAM data set (that is, a data set which shares a data space with other VSAM data sets), VSAM generates a name for these components. The format of the VSAM-generated name is the same as that shown above for a unique data space.

To relate the VSAM-generated name with a VSAM cluster, alternate index, page space, catalog, or data space, you list the catalog that owns the volume. You issue a LISTCAT command to list the catalog's contents. The LISTCAT output listing relates the VSAM-generated names with user-assigned entrynames for cataloged objects.

Each volume owned by a VSAM catalog contains two timestamps that are written in the format-4 DSCB when the volume is first cataloged. Both timestamps are updated but only one (the second) is checked. The first timestamp is maintained for compatibility with earlier VSAM releases. The volume timestamps are updated as follows:

- For a volume owned by a recoverable VSAM catalog, each time the catalog's volume record is modified.
- For a volume owned by a nonrecoverable VSAM catalog, each time space is allocated, extended, or scratched by VSAM. When the volume is mounted, the system compares the second timestamp on the volume to the timestamp in the volume record in the catalog owning the volume. If the volume's timestamp is earlier than the catalog's timestamp, the volume is considered downlevel. Access Method Services will not normally open a data set on a down-level volume.

Some of the implications of VSAM volume ownership are:

- All VSAM clusters, alternate indexes, page spaces, and data spaces on a volume must be cataloged in the catalog that owns the volume. Only one VSAM catalog can own the volume.
- VSAM volume ownership does not affect nonVSAM data sets that reside on the volume. NonVSAM data sets can exist on a volume owned by a VSAM catalog, and can (but should not) be cataloged in a catalog that does not own the volume. Since OS control volumes (CVOLs) are considered nonVSAM data sets, an OS CVOL can also exist on a volume owned by a VSAM catalog.
- In order to release a volume from ownership by a VSAM catalog, you must delete all VSAM objects that reside on the volume. The catalog also contains a volume entry which describes the volume and its VSAM data spaces. After deleting the VSAM objects, you issue the DELETE SPACE command. The DELETE SPACE command deletes the VSAM data spaces on the volume, removes the volume entry from the catalog, and revises the format-1 and format-4 DSCBs in the volume's VTOC. You can delete all VSAM data spaces on a volume whether or not they contain VSAM objects by specifying the FORCE option on the DELETE SPACE command.
- A user catalog cannot exist on a volume that contains active page spaces, paging data sets (that is, system data sets that are paged into and out of the CPU's virtual storage), or the SYS1.STGINDEX. These objects are cataloged in the master catalog and, therefore, the volume(s) on which they reside are owned by the master catalog.

VSAM'S USE OF CATALOGS

Most uses of Access Method Services involve doing something to the VSAM catalogs. For example, establishing a VSAM data set involves creating an entry in a catalog; deleting a VSAM data set involves removing an entry from the catalog; and moving a VSAM data set from one system to another involves moving an entry from a VSAM catalog in one system to a VSAM catalog in another system. The use of Access Method Services with VSAM requires an understanding of how VSAM uses catalogs.

VSAM uses catalogs as a central information point for all VSAM data sets and the direct-access storage volumes on which they are stored. There are two kinds of VSAM catalogs: master catalogs and user catalogs. In addition, OS CVOLs (control volumes) can be attached to a VSAM master catalog. In an OS/VS2 system, the VSAM master catalog is the system's primary catalog and is required. Any number of VSAM user catalogs are optional. The system's nucleus contains a pointer to the VSAM master catalog, which in turn points with catalog connector entries to any VSAM user catalogs. The OS CVOLs are cataloged in the master catalog as nonVSAM data sets with the name "SYSCATLG.catname."

Catalogs provide VSAM with the information to allocate space for data sets, verify authorization to gain access to them, compile usage statistics on them, and relate RBAs to physical locations.

For a VSAM data set to exist, it must be defined in a VSAM catalog. That is, you must enter in the catalog a data set's name and other facts about it by using Access Method Services to define it.

All VSAM data sets on a volume must be cataloged in the same catalog, either the master catalog or a user catalog. A data set is defined in only one catalog. VSAM "ownership" of a volume by a catalog is established the first time VSAM space is allocated on that volume. The DEFINE with the UNIQUE attribute (of a cluster, alternate index, or page space), DEFINE USERCATALOG, or DEFINE SPACE commands establish such ownership. All subsequent VSAM data sets defined on the volume are required to be defined in the same catalog (that is, the catalog that owns the volume).

When you execute a program to process a data set, catalogs are searched to find out which volume(s) the data set is stored on, unless you give volume serial number(s) by way of JCL.

The order in which catalogs are searched depends on whether catalogs are specified for the current job step (STEPCAT) or job (JOB CAT), whether the name of the data set is a qualified entryname, and on whether the catalog is to be used to define a new entry or searched for existing entries. See "Order of Catalog Use" for an explanation of how a catalog is selected during the processing of an Access Method Services command. ("Order of Catalog Use" sections appear in the chapters that describe each of the following commands: DEFINE, BLDINDEX, ALTER, LISTCAT, and DELETE.)

Note: An unqualified name and a qualified name cannot exist in the same catalog if the first qualifier of the qualified name is the same as the unqualified name. For example, DATA and DATA.PAYROLL cannot exist in the same catalog.

Figure 2 illustrates how data sets in an OS/VS2 system can be divided up for cataloging among the VSAM master catalog, VSAM user catalogs, and OS control volume catalogs (CVOLs).

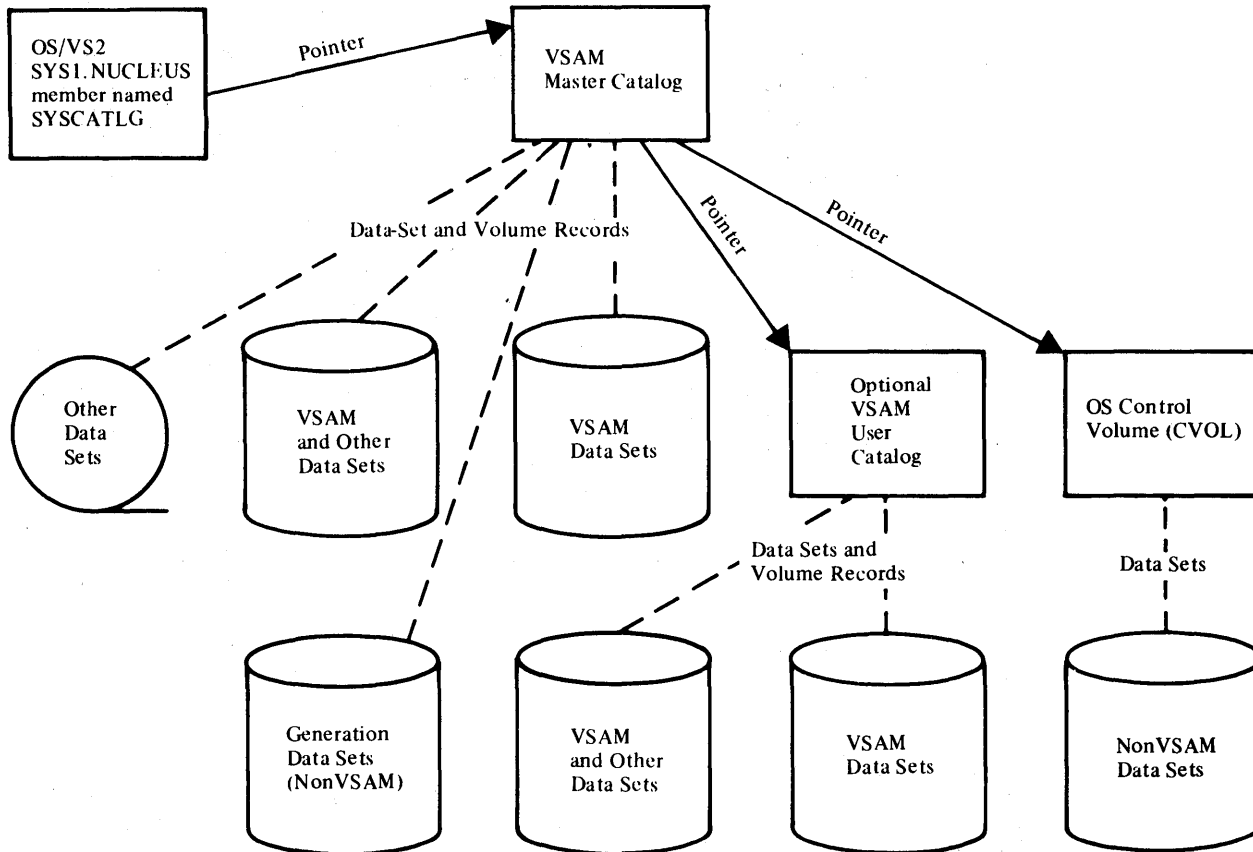


Figure 2. Catalog Relationships in an OS/VS2 System

If you need information on the structure of a VSAM catalog or more details on the information contained in the catalog, refer to Planning for Enhanced VSAM Under OS/VS.

THE MASTER CATALOG

In an OS/VS2 system, the VSAM master catalog is the system's primary catalog. When you generate the system (that is, when you use the SYSGEN procedure), you define the system's master catalog during the first step in Stage II of SYSGEN. The name you specify for the catalog must be different from the name of the master catalog of the driving system. See OS/VS2 System Programming Library: System Generation Reference for details on the system generation process.

The master catalog does not have to reside on the IPL volume. The SYSCATLG member of SYS1.NUCLEUS, a nonVSAM data set on the IPL volume, points to the volume containing the master catalog. The SYSCATLG member is referenced during IPL and NIP processing. It is recommended that the master catalog and VSAM objects cataloged in it not reside on the IPL volume—if the master catalog has to be recovered, then restoring the master catalog will not affect the IPL volume.

The master catalog is the only catalog referenced during NIP processing. Therefore, all system data sets, data sets entered in LNKLSxxx, paging data sets, and libraries specified in the JES start procedure must be cataloged in the master catalog, and not in a user catalog or an OS control volume (CVOL).

You cannot move the master catalog from one system to another by using the EXPORT and IMPORT commands. You can, however, use the REPRO command to make a copy of the master catalog, then move the copy to another system. If your master catalog is recoverable, you can use the LISTCRA, EXPORTRA, IMPORTRA, and RESETCAT commands to repair damaged entries that result from system failure. "Appendix H: Changing the Attributes of the VSAM Master Catalog" describes how you can create a new master catalog for your system. This procedure assumes that you already have a valid master catalog with which to IPL your system. If your master catalog has somehow been destroyed, you can create a new one on another system using the appropriate portion of the same procedure described in Appendix H.

A master catalog cannot be used simultaneously as the VSAM master catalog for two OS/VS2 systems (that is, the master catalog cannot be shared between two systems as each system's master catalog). However, one system's master catalog can be used as a user catalog on another OS/VS2 system that includes VSAM when the catalog is on a shared direct-access device. When a catalog is used simultaneously by two or more systems, all of the catalog's volumes must be on shared direct-access devices. If you do this, take care to assign passwords to each of the catalog's page spaces and system data sets to prevent their accidental or unauthorized use.

TRANSPORTING USER CATALOGS

You can move a user catalog to another system that includes VSAM (that is, to an OS/VS1 or OS/VS2 system or to a DOS/VS system) to make the objects cataloged in the user catalog available to the other system. The user catalog might own only the volume on which it resides. In this case, the volume can be moved as a single unit to another system. When the catalog owns more than one volume, either all volumes owned by the catalog must be moved to the other system or the catalog's ownership of the volumes that are not moved must be taken away before the catalog is moved.

You can use the EXPORT command with the DISCONNECT parameter to remove a user catalog from its OS/VS system. The EXPORT command deletes the user catalog's connector entry in the master catalog. The catalog, its cataloged objects, and its volume(s) become unavailable to the old system.

You use the IMPORT command with the CONNECT option to make the exported catalog available to the new OS/VS1, OS/VS2, or DOS/VS system. The IMPORT command builds a connector entry in the master catalog that identifies the user catalog's volume. If you want the connector entry built to another user catalog's volume, the CATALOG parameter must be specified and must agree with the STEPCAT or JOBCAT DD statement.

JOBCAT AND STEPCAT DD STATEMENTS

User catalogs can be dynamically allocated. User catalogs can also be allocated directly, by using a JOBCAT or STEPCAT DD statement. A previous section, "JCL and Dynamic Allocation," discusses dynamic allocation of catalogs and cataloged objects.

You can specify a catalog that is to be available for use only with a job step, or that is to be available for the entire job. A STEPCAT DD statement identifies a catalog that is available for a single job step; a JOBCAT DD statement identifies a catalog that is available during the entire job. When both JOBCAT and STEPCAT catalogs are specified, the user catalog specified by STEPCAT is used for the job step. The JOBCAT and STEPCAT statements you code should be in the form:

```
//JOBCAT DD DISP={OLD|SHR},  
// DSNAME=usercatname
```

and

```
//STEPCAT DD DISP={OLD|SHR},  
//          DSNNAME=usercatname
```

Because the master catalog contains information about each user catalog, you do not need to specify the catalog's volume and device information.

You can specify more than one user catalog for a job or job step. To indicate another user catalog, follow the JOBCAT or STEPCAT DD statement with a concatenated DD statement—an unlabeled DD statement that specifies the name of another user catalog:

```
//STEPCAT DD DISP=SHR,DSN=MGMTCAT1  
//          DD DISP=SHR,DSN=MGMTCAT2
```

You can supply a JOBCAT or STEPCAT DD statement to allocate the catalog when:

- Your job processes a user catalog (that is, when your sequence of Access Method Services commands adds, modifies, lists, or deletes a VSAM catalog entry)
- Your job includes a space allocation request directed toward a volume that contains VSAM data spaces (that is, a DEFINE SPACE or DELETE SPACE command)
- Your job causes a VSAM data set to be opened and that data set is defined in a user catalog (that is, a REPRO or PRINT of a VSAM data set, a DELETE with the ERASE option, EXPORT, IMPORT, IMPORTRA, BLDINDEX, or VERIFY)

Refer to "Copy-Catalog Procedure" for information on REPRO and JOBCAT/STEPSCAT DD statements.

The description of each of the commands that adds, modifies, lists, or deletes VSAM catalog entries contains a section, "Order of Catalog Use," which explains how VSAM selects a catalog for use by a particular job.

You must supply either a JOBCAT or STEPCAT DD statement to allocate the catalog when:

- Your job uses IEHDASDR to dump a volume owned by a user catalog
- Your job causes the catalog to be opened as a data set (for example, using the PRINT command to print a catalog, using the LISTCRA command with the COMPARE option, using the VERIFY command to restore end-of-file indicators of a user catalog)
- Your job causes the scheduler allocation of the user catalog to be bypassed because unit and volume information were specified on the DD statement describing the VSAM data set defined in that user catalog

You can minimize the use of JOBCAT and STEPCAT DD statements for your jobs when you name your data set with a qualified entryname whose first qualifier is the name or alias of the catalog in which the data set is defined. When the catalog is not identified with a JOBCAT or STEPCAT DD statement or explicitly named (that is, with an Access Method Services command's CATALOG parameter), the OS/VS scheduler searches the master catalog for the data set's entryname. If the entryname is not found, the system uses the entryname's first qualifier as a search argument and attempts to locate either a user-catalog connector entry or a user catalog's alias entry. If the system finds a user-catalog connector entry (that is, an entry whose name or alias is the same as the entryname's first qualifier), the system searches that user catalog for the data set's entry, using the data set's full entryname.

OS CONTROL VOLUMES (CVOLS)

In addition to the VSAM master catalog, your system can include user (or private VSAM) catalogs and OS control volumes (CVOLs or private OS catalogs). The system's catalog structure can be organized with a master catalog only, or with a master catalog and any combination of one or more user catalogs and/or OS CVOLs.

The master catalog (the primary level) points to each user catalog and OS CVOL in the system. A VSAM user catalog cannot be accessed via another user catalog or OS CVOL.

The VSAM master catalog provides the following benefits in your OS/VS2 system:

- Portability of catalogs between systems
- The capability of mounting, as needed, the volumes containing catalogs
- Minimization of the effect of errors or damaged entries on a catalog
- Performance improvements by reducing contention for a catalog

The following differences between VSAM user catalogs and OS CVOLs might affect your choice of having user catalogs or OS CVOLs on your system:

- You cannot catalog VSAM objects in an OS CVOL.
- You cannot use Access Method Services commands to catalog objects in, or modify existing entries of, an OS CVOL.
- You can convert an OS CVOL's entries to VSAM catalog entries when you use the CNVTCAT command, but there is no corresponding facility to convert VSAM catalog entries to OS CVOL entries.
- A user catalog requires at least twice as much space as a corresponding OS CVOL.
- User catalogs can be moved only to other OS/VS systems that include VSAM. OS CVOLs can be moved to MVT and OS/VS2 systems that might not include VSAM.

You can use the DEFINE NONVSAM command to define a CVOL (as though it were a nonVSAM data set) in the master catalog when you name it 'SYSCATLG.qualifier'. You can't use IEHPROGM to create or extend a CVOL.

Many users choose to use OS CVOLs initially, then convert the CVOLs to user catalogs when their system is running smoothly. Whether you use an OS CVOL or a user catalog, nonVSAM data sets are cataloged and uncataloged by your specification of the DD statement's DISP parameter (see OS/VS2 JCL for more details on DD statements and job control language). VSAM objects are cataloged and uncataloged in VSAM catalogs by using Access Method Services commands.

TIME SHARING OPTION (TSO)

TSO is a subsystem of OS/VS2 that provides conversational time sharing from remote terminals. You can use TSO with VSAM and Access Method Services to:

- Execute Access Method Services commands
- Execute a program to call Access Method Services

When TSO is used with Access Method Services, the following differences must be observed:

- TSO allows the initial characters of a keyword to be supplied as an abbreviation of the keyword. The only restriction is that enough initial characters must be supplied to make the keyword unique. TRACKS, for example, could be abbreviated TR, TRA, or TRAC, because no other keyword within the same command can be abbreviated in the same way.

Some abbreviations which are acceptable to Access Method Services such as CYL, CYLINDER, REC, RECORD cannot be used under TSO because the abbreviations do not contain enough initial characters to make the keyword unique. TSO can't tell whether you mean CYLINDERS or CYLINDERFAULT when CYL or CYLINDER is used. Similarly, abbreviations REC and RECORD are ambiguous; TSO doesn't know if you mean RECORDS or RECORDSIZE.

- In addition, the abbreviations described in the remainder of this publication (for example, TRK for TRACKS) are acceptable to TSO.
- When a parameter's value consists of a list of one or more parenthesized parameter sets, the outer parentheses surrounding the list are always required. For example, if lowkey and highkey form a parameter set that can be repeated several times, then the outer parentheses are required even when just one parameter set is specified; as follows:

KEYWORD((lowkey highkey))

- A name can be specified in quotation marks or not in quotation marks. Under TSO, however, a prefix (for example, the userid) is added to a name that is not in quotation marks. The prefix becomes the first qualifier in the name. If the name is a volume serial number (as it may be in the DELETE command, for example), enclose it in quotation marks; if you do not, a prefix is added to the volume serial number.
- Under TSO, a volume serial number can consist of alphabetic, national, numeric, or special (hyphen only) characters; if any other characters are used, the volume serial number cannot be used as an entryname under TSO.
- Under TSO, a password can be supplied with any entryname; the password is ignored if not required.

Note: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password for a cluster. Checking the logon password counts as one attempt to obtain a password. If the user has not specified ATTEMPTS in his DEFINE command, he has one attempt to supply the catalog's password because the default is two.

- The modal commands, used to control execution (IF-THEN-ELSE command sequence, DO-END command sequence, SET, and PARM), are not allowed under TSO.

- Under TSO, the user is prompted to complete a fully-qualified name. The user is also prompted to supply required, but omitted, parameters.
- Under TSO, the SYSPRINT data set is not used. The OUTFILE parameter may be used to specify a data set to receive the output procedure from Access Method Services commands.

For details about the format of a displayed catalog entry (resulting from a LISTCAT command) for a TSO user, see "Appendix B. Interpreting LISTCAT Output Listings."

For details about writing and executing programs and allocating data sets with TSO, see OS/VS2 TSO Terminal User's Guide and OS/VS2 TSO Command Language Reference.

MASS STORAGE SYSTEM (MSS)

The IBM 3850 Mass Storage System can be used with OS/VS2 to store a massive amount of data online to the operating system. It is described in the Introduction to the IBM 3850 Mass Storage System (MSS).

When you have the Mass Storage System, you can define VSAM data spaces, user catalogs, and data sets and nonVSAM data sets on mass-storage volumes. The master catalog cannot be stored on a mass-storage volume.

A VSAM catalog may have defined in it both objects stored on direct-access storage volumes and objects stored on mass-storage volumes. In particular, the data component of a key-sequenced cluster may be stored on a mass-storage volume and the index component on a direct-access storage volume, or vice versa.

Space for an object larger than one cylinder that is stored on a mass-storage volume should be allocated in cylinders to optimize data transfer (staging and destaging) between mass-storage and direct-access storage.

Data stored in the Mass Storage System is staged from mass storage to a direct-access storage staging drive when the object to which the data belongs is opened or when the data is requested. It is destaged from the staging drive to mass-storage when the object is closed. (The Introduction to the IBM 3850 Mass Storage System (MSS) tells what direct-access storage devices can be used for staging.)

Access Method Services for the Mass Storage System provides a set of commands for the management of mass-storage volumes. These commands are described in OS/VS Mass Storage System (MSS) Services: Reference Information.

Access Method Services for managing VSAM catalogs provides parameters for options of the Mass Storage System in the DEFINE and ALTER commands, which enable you to specify how a VSAM data set that is stored on a mass-storage volume is to be staged and destaged and how a user catalog that is stored on a mass-storage volume is to be destaged. These parameters are described in this book.

The staging attributes, BIND, CYLINDERFAULT, and STAGE, affect performance. BIND causes an object to be staged when it is opened and to be retained (bound) in direct-access storage. CYLINDERFAULT causes portions of an object to be staged only as needed during processing. STAGE is a compromise: the object is staged when it is opened, but not retained in direct-access storage. When the staging activity of other objects is light, STAGE can achieve results similar to BIND: the data might remain in direct-access storage, available for requests for access without staging. When the staging activity of other objects is heavy, STAGE can achieve results similar to CYLINDERFAULT: the

data might not remain in direct-access storage and might have to be restaged when needed.

A user catalog that is stored on a mass-storage volume is always staged and bound when it is opened—that is, it is retained in direct-access storage until it is closed. Not binding a user catalog might degrade performance.

A CRA is opened when it is first needed, and will remain open for as long as the user catalog is open. Because of this, the CRA volume must remain mounted during that period of time.

The destaging attributes, DESTAGEWAIT and NODESTAGEWAIT, affect data integrity. DESTAGEWAIT causes VSAM to return control to the program that closes an object synchronously—only after destaging is complete. VSAM can notify the program whether destaging was successful. NODESTAGEWAIT causes VSAM to return control to the program asynchronously—as soon as the object is closed, but before it has been destaged.

A failure in destaging causes a message to be written to the operator. With DESTAGEWAIT, an error code is also returned from CLOSE to the processing program. But there are no recovery procedures for a program to undertake: one use of DESTAGEWAIT is for a processing program to periodically issue a temporary CLOSE of a bound object to cause it to be destaged at various checkpoints. The processing program can terminate processing if a failure occurs in destaging. The last copy destaged would be the copy to fall back to, pending correction of the error that caused the failure.

A data component that is defined with the ERASE parameter and stored on a mass-storage volume, is overwritten with binary zeros on the staging drive after it is destaged.

Note: On MSS volumes, a GDG data set may not be scratched when it is automatically uncataloged, even if the SCRATCH option was specified in the DEFINE.

DATA SECURITY AND PROTECTION

The protection of data includes data security, or the safety of data from theft or intentional destruction, and data integrity, or the safety of data from accidental loss or destruction. Security and integrity options are specified by an entry in the DEFINE command.

DATA-SET SECURITY

Access Method Services provides options to protect data sets against unauthorized use and loss of data. To effectively use the protection features, you must understand the difference between two different operations: 1) referring to a catalog entry and 2) using the data set represented by a catalog entry. A catalog entry is referred to when new entries are defined (DEFINE), or old entries are altered (ALTER), deleted (DELETE), or listed (LISTCAT). The data set represented by the catalog entry is used when it is connected to a user's program (OPEN), or disconnected (CLOSE), or when it reaches its upper RBA-boundary (End-of-Volume). The distinction between these two operations is the key to understanding how the VSAM passwords work; different passwords may be needed for the two operations.

The data-set security options are described in the sections that follow.

PASSWORDS TO AUTHORIZE ACCESS

You can optionally define passwords for clusters, cluster components (data and index), page spaces, alternate indexes, alternate index components (data and index), paths, and VSAM catalogs, which a person must give to get permission to gain access to them. There are different passwords for various degrees of security. The higher levels provide greater protection than the lower levels. The levels are:

- Full access. This is the master password, which allows you to perform all operations (retrieving, updating, inserting, and deleting) on a data set and any index and catalog record associated with it. Using this password to gain access allows you to delete an entire VSAM data set and to alter any catalog information (including passwords) about the data set, index, or catalog, except the cataloged object's physical-extent descriptors. The master password allows all operations and bypasses any additional verification checking by the user-security-verification routine.
- Control access. This password authorizes you to use control-interval access. See OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications for information on the use of control-interval access.
- Update access. This password authorizes you to retrieve, update, insert, or delete records in a data set. Specifying a catalog's update password authorizes you to define data sets in it, and to alter or delete nonVSAM and GDG entries in it.
- Read access. This is the read-only password, which allows you to examine data records and catalog records, but not to add, alter, or delete them, nor to see password information in a catalog record.

If you define passwords for any data sets in a catalog, you must also define passwords for the catalog in order for the data-set passwords to have effect. If you do not define passwords for the catalog, no password checking will take place during operations on the data set's catalog entries.

Operations on a catalog may be authorized by the catalog's appropriate password or, in some cases, by the appropriate password of the data set whose definition in the catalog is being operated on. For example, defining a data set in a password-protected catalog requires the catalog's update (or higher) password. Listing or deleting a definition requires the appropriate password of either the catalog or the data set. However, if the catalog, but not the data set, is protected, you need give no password to list the data set's catalog definition.

Because a user catalog defines itself, it may be password-protected without the master-catalog being password-protected. To delete a user catalog, you must give its master password, whether the master catalog is password-protected or not.

Each higher-level password allows all operations permitted by lower levels. Any level may be null (not specified), but if a low-level password is specified, the master-level password must also exist. The DEFINE and ALTER commands accomplish this by giving all of the higher passwords the value of the highest password specified. Thus, if you specify only a read-level password, that password will become the update-, control-, and master-level password as well. If you specify a read password and a control password, the control-password value will become the master-level password as well. However, the update-level password will be null.

One reason for password-protecting the components of a cluster is to prevent access to the index of a key-sequenced data set, since the only way to gain access to an index is to open it independently of the cluster. (See OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications for a description of access to an index.)

Some Access Method Services operations may involve more than one password authorization. For example, importing a data set involves defining the data set and loading records into it. If the catalog into which the data set is being imported is password-protected, its update- (or higher-level) password is required for the definition; if the data set is password-protected, its update- (or higher-level) password is required for the load. The IMPORT and IMPORTRA commands allow you to specify the password of the catalog; the password, if any, of the data set being imported is obtained by the commands from the exported data.

Every VSAM data set is represented in the VSAM catalog by two or more entries: a cluster entry and a data entry, or, if the data set is a key-sequenced data set, a cluster entry, a data entry, and an index entry. Of the two or three entries, the cluster entry is the controlling entry. Each of the two or three entries can have its own set of four passwords; the passwords you assign have no relationship to each other. For example, if you password-protect a cluster but don't password-protect the cluster's data component, someone could issue LISTCAT to determine the name of your cluster's data component, then open the data component and access records in it—even though the cluster itself is password protected.

Catalogs are themselves VSAM data sets, and may have passwords. For some operations (for example, listing all of a catalog's entries with their passwords or deleting catalog entries), the catalog's passwords may be used instead of the entry's passwords. Thus, if the master catalog is protected, the update- or higher-level password is required when defining a user catalog. This is because all user catalogs have an entry in the master catalog. When deleting a protected user catalog, the user

catalog's master password must be specified.

The following protection considerations and precautions should be observed when using commands which refer to the VSAM catalog:

- Passwords are 1 to 8 EBCDIC characters, either alphameric or special characters, and may use hexadecimal format (see the CONTROLPW, MASTERPW, READPW, and WRITEPW definitions in the "Command Format" chapter).
- To create a catalog entry (with the DEFINE command), the update- or higher-level password of the catalog is required.
- To modify a catalog entry (with the ALTER command), the master password of the entry or the master password of the catalog which contains the entry is required. However, if the entry to be modified is a nonVSAM or GDG entry, the update-level password of the catalog is sufficient.
- To gain access to passwords in a catalog (for example, to list or change passwords), you must specify the master-level password of either the entry or the catalog. Similarly, a master-level password must be specified with the DEFINE command if you want to model an entry's passwords.
- To delete a protected data-set entry from a VSAM catalog requires the master-level password of the entry or the master-level password of the catalog containing the entry. However, if the entry describes a VSAM data space, the update-level password of the catalog is sufficient. If the entry to be deleted is a nonVSAM, GDG, or alias entry, the update level password of the catalog is sufficient.
- If the password of the catalog is the read-level, catalog entries with the read-level passwords may be listed by specifying the read password of the entry or the catalog's read level password. However, entries without passwords may be listed without specifying the catalog's read-level password. If you wish to list the passwords associated with a catalog entry, you must specify the master password of the entry or the catalog's master password.
- If the proper password is not specified with the Access Method Services command, a password prompt will occur. Unless you have specified the CODE parameter on either the DEFINE or ALTER command, the prompt will include the name of the data set; if you have specified CODE, the prompt will include the code name you specified. In some circumstances, due to the manner in which catalog entries are referenced, more than one prompt may occur. For example, when an ALTER or DELETE request is processed, the catalog must be referred to twice, once to locate the information, and once to perform the requested function. Again, incorrect password specification when listing catalog entries may cause numerous prompts. To avoid unnecessary prompts, specify the catalog's password, which allows access to all entries that the operation affects. A catalog's master-level password will allow you to refer to all catalog entries. However, a protected cluster cannot be processed with the catalog's master password.
- Specification of a password where none is required is always ignored.

The following protection considerations and precautions should be observed when using commands which cause a data set access:

- Passwords are 1 to 8 EBCDIC characters, either alphameric or special characters, and may use hexadecimal format (see the CONTROLPW, MASTERPW, READPW, and WRITEPW definitions in the "Command Format" chapter).

- To access a VSAM data set using its cluster name, instead of data or index names, you must specify the proper level password for the cluster. The proper level password for the cluster is required even if the data or index passwords are null.
- To access a VSAM data set using its data or index name, instead of its cluster name, you must specify the proper data or index password. However, if cluster passwords are defined, the master password of the cluster may be specified instead of the proper data or index password.
- If a cluster has only null (not specified) passwords, you can access the data set using the cluster name without specifying passwords. This is true even if the data and index entries of the cluster have passwords defined. This allows unrestricted access to the VSAM data set as a whole but protects against unauthorized modification of the data or index as separate components.

Operator Prompting Code

Computer operators and TSO-terminal users may be given the opportunity to supply a correct password if a processing program does not give the correct one when it tries to open a password-protected data set. When the data set is defined, you may specify a code to be used in lieu of the data-set name to prompt the operator or terminal user for a password. The prompting code keeps your data secure by not allowing the operator or terminal user to know both the name of the data set and its password.

A data set's code is used for prompting for any operation against a password-protected data set. The catalog code is used for prompting when the catalog is opened as a data set, when an attempt is made to locate catalog entries that describe the catalog, and when an entry is to be defined in the catalog.

If you don't specify a prompting code, VSAM identifies the job for which a password is needed with the JOBNAME and DSNAME for background jobs or with the DSNAME alone for foreground (TSO) jobs.

Attempts to Supply a Password

When you define a data set, you can specify the number of times the computer operator or terminal user is allowed to try to give the correct password when a processing program is trying to open a data set. If the allowed number of attempts is exceeded and you are using the System Management Facilities, a record is written to the SMF data set to indicate a security violation.

Note: Using the TSO logon password counts as one attempt.

PASSWORDS FOR NONVSAM DATA SETS

When you define a nonVSAM data set in a VSAM catalog, the data set is not protected with passwords in its catalog entry. You can password-protect a nonVSAM data set when you create it, by specifying LABEL=(PASSWORD|NOPWREAD) in the DD statement that describes the data set (for more details, see OS/VS2 JCL). You use the PROTECT macro instruction to assign a password to the nonVSAM data set (for more details, see OS/VS2 Data Management Services Guide and OS/VS2 System Programming Library: Data Management).

If the catalog is update protected, you must supply the catalog's update- (or higher-level) password in order to define, delete or alter a nonVSAM data set. The password can be supplied as a subparameter of the command's CATALOG parameter, or as a response to the password-prompting message.

USER-SECURITY-VERIFICATION ROUTINE

In addition to password protection, VSAM allows you to protect data by specifying a program to verify a user's authorization. Specific requirements of the user-security-verification routine are described in the chapter "User-Written Exit Routines" in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide. To use this routine, simply specify the name of the authorization routine you have written in the AUTHORIZATION parameter of the DEFINE or ALTER command.

If a password exists for the type of operation being performed, the password must be given, either in the command or in response to prompting. The user-security-verification routine is called only after the password specified is verified. The user-security-verification routine is bypassed whenever a correct master password is specified, whether or not the master password is required for the requested operation.

RESOURCE ACCESS CONTROL FACILITY (RACF) PROTECTION

Resource Access Control Facility (RACF) provides an optional software access control measure you can use in addition to or instead of passwords. Password protection and RACF protection can coexist for the same data set. When RACF protection is applied to a data set which is already password protected, password protection is bypassed and access is controlled solely through the RACF authorization mechanism. If a User-Security-Verification Routine (USVR) exists, it will not be invoked for RACF-defined data sets.

Although passwords are ignored for a RACF indicated data set, they can still provide protection if the data set is moved to another system which does not have RACF protection.

RACF authorization checking is generally compatible with the password authorization checking scheme. The compatibility includes the time at which the authorization check is made and the sources of authorization. The RACF authorization levels of alter, control, update, and read correspond to the password levels of master, control, update, and read in VSAM. As an example, consider the authorization required to delete a VSAM data set. With passwords, this requires master level authorization either to the data set or to the containing catalog. (The catalog is a second source of authorization in this case.) If both the data set and catalog were RACF protected, it would take alter level authorization for the catalog or data set to delete the data set.

One difference between password authorization checking and RACF authorization checking is that deletion of any type of RACF-protected entry from a RACF-protected VSAM catalog requires alter level authorization either to the catalog or to the entry being deleted.

In order for RACF protection to be applied to a data set, the containing catalog must also be RACF protected. In order for password protection to be applied to a data set, the containing catalog must be either RACF protected or password protected and the data set itself not defined to RACF.

Refer to OS/VS2 MVS Resource Access Control Facility (RACF) General Information Manual for more information.

PROTECTING USER'S DATA

Sometimes, because of system or hardware failure, your data may be destroyed or made inaccessible to you. You should recognize the possibility of this occurring and take actions that will allow you to recover your data if such a catastrophe happens.

PREFORMATTING CONTROL AREAS

When you define a cluster, you can indicate that VSAM is to preformat each control area as records are loaded into the cluster (RECOVERY) or is not to preformat them, in the interest of performance (SPEED). Preformatting clears all previous information from the direct-access storage area and writes an end-of-file indicator:

- For an entry-sequenced or relative record data set, in every control interval in the control area.
- For a key-sequenced data set, in the first control interval in the control area following the preformatted control area. (The preformatted control area contains free control intervals.)

As records are loaded into a preformatted control area, there is always a following end-of-file indicator that indicates how far loading has progressed. If an error occurs that prevents loading from continuing, you can readily identify the last successfully loaded record and resume loading at that point.

Without preformatting, an end-of-file indicator is written only after the last record is loaded. If an error occurs that prevents loading from continuing, you might not be able to identify the last successfully-loaded record—you might have to reload the records from the beginning.

BACKING UP DATA

Two Access Method Services commands enable you to create a backup copy of your data:

- EXPORT copies a VSAM cluster or alternate index and its catalog entries. The copy can be imported to another OS/VS1, OS/VS2, or DOS/VS system that includes VSAM or can be used as a backup.
- REPRO copies a VSAM cluster or alternate index (without its catalog entry), a VSAM catalog, or a sequential data set. The copy can be used as a backup.

The choice of using EXPORT or REPRO as a backup depends on a number of factors:

- A backup copy created by EXPORT can only be a sequential data set. A backup copy created by REPRO can be a VSAM data set or a sequential data set. Also, the backup copy obtained by using REPRO can be a different type of VSAM data set than the original. You could, for example, back up a VSAM key-sequenced data set by copying it to a VSAM entry-sequenced data set using REPRO. By using the REUSE attribute when using REPRO to back up a VSAM data set, you have the facility of taking repeated backups to the same VSAM reusable target data set.
- A backup copy created by EXPORT contains all the information necessary to redefine the VSAM cluster or alternate index when you IMPORT the copy. A backup copy created by REPRO does not contain this catalog information. If it is necessary to use the backup copy, you must delete the original data set, if it still exists, using the DELETE command. You must then redefine the data set using the

DEFINE command prior to restoring it with the backup copy using the REPRO command. (Note that the delete and define operations would not be necessary if the original data set still exists and was defined with the REUSE attribute.)

Using EXPORT and IMPORT to transport or back up a key-sequenced data set or using REPRO to copy it reorganizes the data set—it rearranges data records physically in ascending key sequence (control-interval and control-area splits might have got them physically out of order) and balances free space quantities.

After you have replaced a damaged data set with its backup copy, you can bring the backup copy up to date with the original by rerunning the jobs that updated the original between the time it was backed up and the time it became inaccessible.

Backing up the data sets in a user catalog enables you to recover from damage to the catalog. You can import the backup copy of a data set whose entry was lost or redefine the entry and reload the backup copy. If the catalog was completely lost, you could redefine it, then import or redefine and reload all of the data sets that were defined in the catalog. (Backing up the catalog itself is discussed in the section "Protecting VSAM Catalogs.")

Exporting and Importing a Data Set

You can use the EXPORT command (with the TEMPORARY parameter) to copy a VSAM cluster and its catalog entries onto a movable volume (that is, a demountable direct-access volume or a magnetic-tape volume). You can then use the IMPORT command either to move the portable copy to another VS1, VS2, or DOS/VS system that includes VSAM or to replace the original cluster and its catalog entries. The portable copy itself is inaccessible for processing.

Making a Copy of a Data Set

You can use the REPRO command to copy a VSAM or sequential data set into another VSAM or sequential data set, newly defined (for protection) in another catalog. (The REPRO command does not copy the data set's catalog information.) You can use the REPRO command with the REPLACE parameter to merge the backup copy into the original data set, or you can delete and redefine the original data set and use REPRO to reload the backup copy into it. Or, because the backup copy is itself accessible for processing, you can replace the original with it. (You have to make available for processing the catalog in which you defined the backup, and you have to use the name of the backup.)

If you periodically process a data set sequentially, you can easily create a backup copy as a by-product of normal processing. You can use this backup copy like one made by way of REPRO.

ENSURING ACCESSIBILITY OF SECONDARY EXTENTS

When a VSAM catalog is damaged so that you have to use a backup copy of the catalog (see "Protecting VSAM Catalogs," below, for information about backing-up catalogs), any extents that were allocated to the catalog's data sets since the backup copy was made become inaccessible.

In an entry-sequenced or a relative-record data set, the records in the new extents will have been added at the end. But in a key-sequenced data set, some of the records in the new extents could have been moved there because of a control-area split.

To ensure that all of the records in a catalog's data set are accessible after you fall back to a backup copy of the catalog, you can:

- Eliminate the need for secondary extents in the data set by providing a sufficiently large primary space allocation.
- Reduce the number of secondary extents by providing for sufficiently large secondary space allocations. Whenever a secondary extent is allocated, you can make a new backup copy of the catalog. You can monitor the growth of a data set by way of the LISTCAT command or the SHOWCB macro. (The use of SHOWCB to display statistics is described in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.)

For growth that did not require secondary allocation in a data set since the backup copy of its catalog was made, you can use the VERIFY command to reset the end-of-file indicators in the catalog after you fall back to the backup copy.

CORRECTIVE MEASURES

When you discover that a catastrophe destroyed part or all of your on-line data, you can replace the destroyed data set with its backup copy (if one exists):

- You can use the REPRO command with the REPLACE option to merge the backup copy with the destroyed data. When only part of your data set is damaged, you can specify REPRO parameters that allow you to replace only the records in the damaged part of the data set.
- You can use the IMPORT command to totally replace a VSAM cluster whose backup copy was built using the EXPORT command. The IMPORT command uses the backup copy to replace the cluster's contents and catalog information.

PROTECTING VSAM CATALOGS

Because you cannot gain access to a VSAM data set without the master or user catalog in which the data set is defined, you need to make sure that your catalogs can always be made available. You can take preventive measures to minimize the possibility of catalog loss and corrective measures to recover from loss, should it occur.

The effect of the loss of a catalog can be minimized by the following:

- Widespread use of user catalogs—a user catalog for each volume, where possible. (It is not possible for multivolume data sets—one catalog must control all the volumes of a data set.) Compare the effect of the loss of a catalog when (a) 10 data sets are cataloged in each of 10 catalogs and (b) 50 data sets are cataloged in each of 2. The fewer the catalogs the greater the disruption in the event of loss of a catalog.
- Limited use of the master catalog.
- Dedicated use of volumes for VSAM data sets to segregate VSAM and nonVSAM recovery. You can dedicate a volume by defining a VSAM data space that occupies the whole volume. If some of the space on a volume does not belong to a VSAM data space, OS/VS may allocate that space to a nonVSAM data set. If you follow a dump/restore recovery procedure (outlined later) you will have to fall back to a previous copy not only of the VSAM data and catalog, but also of the nonVSAM data.

USING REPRO TO BACK UP A CATALOG

You can use the REPRO command to unload (make a backup copy of) a catalog. If the catalog becomes inaccessible, you can use REPRO to reload the copy. The section "Backing up a Catalog" in the chapter "Copying and Printing" describes catalog unload/reload.

The backup copy might not reflect current information for some of the cataloged objects. The section "Updating a Backup Catalog" below describes how to make a backup catalog current.

BACKING UP THE MASTER CATALOG

An OS/VS2 system requires a master catalog—if a system or hardware failure damages the master catalog, the system cannot be used until the damage is corrected, unless you have a backup master catalog. Without a copy of the master catalog, you might have to generate a new system (using SYSGEN processing).

Because the system requires a master catalog, the reload procedure described in "Backing up a Catalog" in the chapter "Copying and Printing" requires modification for reloading the master catalog. Reloading the master catalog with REPRO works well, if you have a master catalog to reload and it was not defined as a recoverable catalog. For catalogs defined as recoverable, see "Restoring Catalog Entries After System Failure." When only some of the entries in the master catalog have become inaccessible, and the master catalog itself is still operational, you might be able to reload the backup catalog into it. Otherwise, you can get a master catalog by:

- Using IEHDASDR (on another system) to restore the volume that contains the master catalog (from a tape onto which you have previously dumped the volume). You can do an IPL (initial program load) to bring your system up with the restored volume and then use REPRO to reload the backup into the restored version of the catalog. (If the backup is no more recent than the restored catalog, you can use the restored version without reloading.)

With DUMP/RESTORE, you should design the volume that contains the master catalog to contain information that changes very little. Thus, you won't lose changes when you restore the volume. And, to avoid needing another system to restore the volume from a tape, you could have two direct-access volumes (with the same volume serial number) that contain a copy of the master catalog.

- Using DEFINE (on another VS2 system) to define a user catalog with the same name as your master catalog. You then use REPRO (still on the other system) to reload the user catalog with the backup copy of the master catalog. You can bring up your system with the reloaded user catalog as your master catalog.

DUMPING A CATALOG AND ITS DATA SETS

This procedure calls for periodically dumping the volume(s) upon which a catalog and its data sets are stored with the IEHDASDR utility program (or with some other utility that achieves the same effect). Then, if the catalog is lost or somehow becomes inaccessible, you would be able to restore the volumes (or alternate volumes with the same volume serial numbers). The volumes would then contain what they contained when the backup copies were made. See "Updating a Backup Catalog" below about how to bring the restored catalog up to date.

When you use IEHDASDR to dump a volume containing a VSAM catalog, you should ensure that all update activity (for example, DEFINE, DELETE, ALTER, data set extension, etc.) is quiesced. If any catalog update activity occurs during the dump

operation, it is possible for the chain of free records in the catalog's low key range to be damaged in the backup copy. Then, if it is necessary to restore the backup copy, the damaged chain will be introduced into the restored copy.

If some of the space on a volume does not belong to a VSAM data space, the system may have allocated that space to nonVSAM data sets. With DUMP/RESTORE, you will have to fall back to a previous copy not only of the VSAM data sets and catalog, but also of the nonVSAM data sets.

When you use IEHDASDR to dump or restore a volume that contains VSAM objects (that is, clusters, catalogs, data spaces, and alternate indexes), you must include a JOBCAT or STEPCAT DD statement with the IEHDASDR job or jobstep to describe the VSAM catalog that owns the volume. The volume's table of contents (VTOC) indicates that part of the volume's contents is password-protected, and that the volume is owned by a VSAM catalog (that is, the passwords for the password-protected objects are part of the object's catalog information). However, the volume's VTOC does not identify the catalog that contains the passwords. Before a volume is dumped, the system asks the operator for the correct password of each password-protected object on the volume. However, if the volume is RACF-protected, RACF read authorization to the volume overrides VSAM password protection. For more information on how to use IEHDASDR to dump and restore volumes containing VSAM objects, see OS/VS2 MVS Utilities.

UPDATING A BACKUP CATALOG

A reloaded or restored backup catalog won't reflect any changes in the original catalog between the time you unloaded it to make the backup and the time it became inaccessible. Three types of changes may have occurred:

- Entries may have been deleted by way of DELETE or permanently exported by way of EXPORT PERMANENT. You can remove these entries from the backup catalog with DELETE NOSCRATCH, which is described below under "VSAM Catalog Cleanup."
- Entries may have been defined by way of DEFINE, imported by way of IMPORT, or defined (nonVSAM entries only) by the VS2 Scheduler. You can use DEFINE NONVSAM to bring the backup catalog up to date on nonVSAM entries. But a VSAM object defined by an entry missing from the backup catalog is no longer accessible. To regain the use of the space on a volume or in a VSAM data space whose entry is missing from the backup catalog, you can use ALTER REMOVEVOLUMES (described below under "VSAM Volume Cleanup").

To recover objects that are lost in this way, you will have to redefine them and rerun the jobs that loaded and updated them, or you will have to have backed them up as described above under "Backing Up Data."

- Entries may have been updated to describe extensions. For an extension that used space already allocated to an object (that is, before the original catalog was unloaded to make the backup), you can use the VERIFY command to update the object's end-of-file indicator. But for an extension that used newly-allocated space, you cannot update the backup catalog to recover the data in the extension. (See "Ensuring Accessibility of Secondary Extents" above about how to protect yourself from this loss of data.)

VSAM VOLUME CLEANUP

You normally take away the ownership of a volume from a VSAM user catalog by deleting all the objects and all the data spaces on the volume. But if a user catalog is inaccessible for some reason, or it no longer contains entries for the volume or its data spaces (as, for example, when you reload a backup catalog), you cannot use the DELETE command to take away ownership.

The ALTER command with the REMOVEVOLUMES parameter enables you to get rid of all the VSAM data spaces on a volume without gaining access to the catalog that owns the volume. ALTER REMOVEVOLUMES overwrites the data spaces with binary zeros and rewrites the VTOC to remove the data spaces' format-1 DSCBs and to turn off the VSAM ownership bit in the format-4 DSCB.

When the user catalog itself is on the volume, ALTER REMOVEVOLUMES overwrites it as well. Thus, you can use ALTER REMOVEVOLUMES to clean up a volume that contains a user catalog that has become inaccessible. You should first use EXPORT DISCONNECT to remove the catalog's entry in the master catalog prior to using ALTER REMOVEVOLUMES.

ALTER REMOVEVOLUMES does not remove nonVSAM data sets from a volume, nor does it remove VSAM objects from a volume owned by the master catalog. See the section "VSAM Volume Recovery Function" for details on how to specify the ALTER REMOVEVOLUMES function.

Note:

- ALTER REMOVEVOLUMES is also used simply to take away from a catalog ownership of a candidate volume that does not yet contain a VSAM data space and is not referenced by any VSAM objects.
- You should not use ALTER REMOVEVOLUMES to get rid of VSAM objects on a volume owned by a catalog that you can still use—you should use the DELETE command.
- If you are using the ALTER REMOVEVOLUMES function for a volume containing a user catalog, you should make sure that the user catalog is not open. You can do this by including a DD statement (which is not a JOBCAT or STEPCAT DD statement) which identifies the user catalog and specifies DISP=OLD in the job step that includes the ALTER REMOVEVOLUMES command. The user catalog should not be identified in a JOBCAT DD statement in the ALTER REMOVEVOLUMES job, nor in a STEPCAT DD statement in the ALTER REMOVEVOLUMES job step, nor in any other DD statement(s) or Access Method Services command parameter(s) that would cause the catalog to be allocated and opened as a catalog prior to the time the ALTER REMOVEVOLUMES command is executed (for example, use of a data-set name with a high-level qualifier which is the alias of the user catalog).

If the user catalog that owns the volume is available, you can use the DELETE SPACE command with the FORCE option to get rid of all VSAM data spaces on a volume whether or not they contain VSAM data sets. DELETE SPACE FORCE deletes all VSAM data spaces, marks any VSAM data sets contained therein as unusable in the catalog, and releases ownership of the volume. If your volume contains VSAM data space(s) you wish to eliminate but some VSAM data sets you wish to save, you can use the EXPORT command to export the data sets you wish to save. Then use DELETE SPACE FORCE to eliminate all VSAM data spaces on the volume. After redefining the data spaces, you can reestablish your data sets by using the IMPORT command.

DELETE SPACE FORCE cannot be used for a volume that contains a user catalog. DELETE USERCATALOG FORCE can be used to DELETE a user catalog without first deleting all its entries. All data

spaces are deleted from each volume owned by the catalog, including the catalog volume itself, and volume ownership is released by VSAM.

CHANGING A VOLUME'S SERIAL NUMBER

When you change the volume serial number of a volume that contains VSAM objects, the VSAM objects can no longer be located when the objects' catalog volume information is referenced (you can use the LABEL statement of the IEHDASDR program to change a volume's serial number, or "clip" the volume). You should use the following procedure to change the volume's serial number, to allow the VSAM objects on the volume to be located:

1. Issue the EXPORT PERMANENT command for each VSAM object (that is, for each alternate index and cluster) on the volume. When the volume also contains a VSAM catalog, issue an EXPORT DISCONNECT command to disconnect the user catalog from the master catalog. The EXPORT command copies each VSAM object and its catalog entry onto a movable volume. The object's entry in the user catalog is deleted.
2. Issue the DELETE command to delete all empty data spaces on the volume, and to delete the volume entry from the catalog. NonVSAM data sets are described in the volume's VTOC and are not affected by the DELETE command.
3. When the volume contains an empty catalog (that is, the catalog only describes space on the volume), delete the catalog. If the empty catalog describes space on more than its own volume, you can delete the space on each volume first, then delete the catalog.
4. Execute the IEHDASDR program with the LABEL statement to change the volume's serial number.
5. If the catalog was deleted in step 3, it must be redefined when beginning this step. Issue the DEFINE SPACE command to build a data space on the volume, to establish the VSAM catalog's ownership of the volume, and to build a volume entry that points to the volume with its new serial number. The data space you define should be large enough to contain all of the suballocated VSAM objects you removed from the volume during step 1. (When you execute the IEHDASDR program, you might reorganize the nonVSAM data sets remaining on the volume, so that the space available for the VSAM data space is contiguous.)
6. Issue the IMPORT command for each VSAM object that was removed from the volume during step 1, specifying the new volume information for the object.

VSAM CATALOG CLEANUP

When a volume becomes inaccessible, you can use the DELETE command with the NOSCRATCH parameter to remove the entries for the VSAM objects on the volume from the catalog that owns the volume. You can also use DELETE NOSCRATCH to remove from a backup catalog entries for the VSAM objects that were deleted from the original catalog after the time it was unloaded to make the backup. DELETE NOSCRATCH removes an entry without gaining access to the volume indicated in the entry. You should keep in mind that, if the cluster, alternate index, or pagespace was defined with the UNIQUE attribute, a format-1 DSCB was written on the object's volume. If you specify NOSCRATCH and the format-1 DSCB still exists, it will not be erased from the VTOC. (DELETE without NOSCRATCH specified allows Access Method Services to gain access to a volume to change the VTOC or overwrite an object with binary zeros.)

VSAM cleans up the catalog:

- When DELETE CLUSTER, DELETE ALTERNATEINDEX, or DELETE PAGESPACE is specified, by deleting the cluster, alternate index, or pagespace entry. (Note: NOSCRATCH cannot be specified for these entry types if they are defined in a recoverable catalog.)
- When DELETE SPACE is specified, by deleting the volume entry. The volume entry must be empty (that is, you must first delete the entries of all VSAM objects that are indicated in the entry).

EXPORTING A CATALOG'S DATA SETS

While the primary purpose of the EXPORT command is to allow you to move data sets and user catalogs between systems, it can also be used to create backup copies of particularly sensitive data sets. The EXPORT command copies a data set and extracts the data set's definition from the catalog in which it is defined. If you use Access Method Services to periodically export a data set, you can redefine the catalog should it be lost, define a VSAM data space if the data set is to be suballocated, and import the data set into the new catalog. With export/import you can control the level of each data set individually. If you exported a data set for backup after each run, you wouldn't have to rerun any jobs to bring the backup copy up to date.

A single volume for each catalog is desirable when your system contains a large number of single-volume data sets.

MOVING A CATALOG TO ANOTHER SYSTEM

You can move a user catalog from your OS/VS2 MVS system to another system (that is, to an OS/VS or DOS/VS system that includes VSAM) to make the objects cataloged in the user catalog available to programs running on the other system. The user catalog might own only the volume on which it resides. In this case, the volume, its catalog, and its contents (that is, VSAM clusters, nonVSAM data sets, and free space) can be moved as a single unit to another system. When the catalog owns more than one volume, each volume owned by the catalog must either be moved to the other system with the catalog or be deleted before the catalog is moved.

Caution: When you move a recoverable catalog to a system that does not support recoverable catalogs, the recoverability of the catalog is destroyed when any of the catalog's records are modified, deleted, or added.

You use the EXPORT command with the DISCONNECT option to remove a user catalog from its OS/VS system. The EXPORT command deletes the user catalog's connector entry in the system's master catalog. The EXPORT command also deletes all alias entries that point to the user catalog's connector entry. The catalog, its cataloged objects, and its volume(s) are now unavailable to users of the old system.

When the user catalog has aliases, you should take care to list the aliases before you issue the EXPORT command. You can use the LISTCAT command to list the catalog's aliases, which are in the master catalog. You can use the list of aliases to help you reconstruct the alias entries in the new system's master catalog.

You use the IMPORT command with the CONNECT option to make the exported catalog available to the new OS/VS system. The IMPORT command builds in the new system's master catalog a user-catalog connector entry that points to the user catalog's volume. You should issue a DEFINE ALIAS command to build an alias entry in the new system's master catalog for each of the user catalog's aliases.

CATALOG RECOVERY

All VSAM catalogs can be defined with the RECOVERABLE attribute that makes it possible to recover VSAM data sets and their catalog entries should the catalog be damaged or destroyed. Recovery is achieved by recording a copy of all catalog information about a given volume and its data sets on that volume as well as in the catalog itself. Recovery information is recorded on each volume owned by a catalog. Space for this information is automatically set aside when you acquire VSAM ownership of a volume and also when you define the catalog itself. There is no separate catalog entry for the recovery space: VSAM records its physical track address in the volume's format-4 DSCB.

The recovery information in the volume's catalog recovery area (CRA) is updated immediately whenever parallel information in the catalog is changed. The affected volume(s) must be mounted, and the kind of operation to be performed on an object (data space, cluster, path, etc.) determines which volume(s) to mount.

To recover a VSAM data set and its catalog entries, you issue the EXPORTRA command. EXPORTRA uses the information in the CRA rather than the catalog to gain access to VSAM data sets and produce a copy of the VSAM data sets. The copy can be introduced back into the system by means of the IMPORTRA command.

RESETCAT is another catalog-recovery tool. If, for any reason, inconsistencies develop between your catalog and the CRAs of its owned volumes, you should consider RESETCAT as your recovery vehicle. RESETCAT is basically a synchronization process between a catalog and its associated CRA(s). If your catalog or any of its owned volumes becomes unusable, you can restore a backup volume. However, this will not solve your problems because inconsistencies will undoubtedly exist between your catalog and the CRAs of the volumes it owns. Either the catalog or the CRAs will be down level (entries will not precisely describe the data sets). RESETCAT provides the necessary synchronization facility to ensure consistency between the catalog and its volumes. RESETCAT confines its processing to the catalog and CRAs; VSAM data sets are unaffected by this operation; no data movement takes place.

You can use the LISTCRA command either to list the contents of the CRA before you do selective recovery or to list the entries in the recovery area that are different from those in its associated catalog.

See the "Restoring Catalog Entries After System Failure" chapter for more information and examples on catalog recovery.

If you are presently using an OS/VS2 system without recoverable catalogs, see Appendix H for information on making your master catalog recoverable.

Regaining Access to Data

Using some of the corrective measures listed, you can analyze and recover from the following conditions:

- Inaccessible data set
- Unusable catalog
- Inaccessible volume

You can use the LISTCRA command with the COMPARE option to identify mismatches between the catalog recovery area (CRA) and the catalog itself. The records in the CRA are used as a base for comparison and, whenever a mismatch is detected, LISTCRA prints the CRA record and its matching catalog record and identifies the mismatch by printing asterisks below the mismatched area(s) of the records. (The LISTCRA command, as well

as EXPORTRA, IMPORTRA, and RESETCAT, is usable only with recoverable catalogs.) The mismatches detected by LISTCRA vary in their degree of seriousness. The type of mismatch determines the level of recovery that is required. The following list (ordered by severity) shows the type of mismatch, why it may have occurred, and how serious it is. The "Type" column indicates the section of the catalog record or group of fields in which the mismatch occurred. If the LISTCRA command detects more than one area of mismatch, only the most serious mismatch is identified. Subsequent sections of this chapter tell how to use corrective measures to recover.

Figure 28 in Appendix D shows an example of the LISTCRA output using the DUMP COMPARE options.

Catalog Volume Records

*MISCOMPARES Message	Type	Cause	Severity
DATA SPACE EXTENTS	Mismatched data space group	A difference in the number, size, and/or location of VSAM data spaces.	Requires recovery of the entire volume.
		A difference in the number and/or location of extents for one or more data sets.	Requires recovery of the entire volume.
		Space has been extended or deleted.	Requires recovery of the entire volume.
DATA SET DIRECTORY	Mismatched data set directory	A difference in the names and/or number of data sets associated with this volume.	Requires recovery of the entire volume.

VSAM Object Records

*MISCOMPARES Message	Type	Cause	Severity
CATALOG ENTRY HAS DIFFERENT NAME	Mismatched entryname	The catalog was restored, or the volume containing the data set was restored. As a result the record in the catalog pointed to by the CRA record is no longer for the same object.	Requires data set recovery.
VOLUME OR KEYRANGE	Mismatched volume or key range in the volume information set of fields.	The catalog was restored, or the volume containing the data set was restored. As a result the object's volume locations or keyranges in the CRA do not match those in the catalog.	Requires data set recovery.
EXTENTS	Mismatched extents in the volume information set of fields.	Data set was not properly closed, the catalog was restored, or the volume containing the data set was restored.	Requires data set recovery.
HIGH USED RBA	Mismatched high used Relative Byte Address	Same as for mismatched extents.	Requires use of the VERIFY command to correct the high RBA.
STATISTICS	Mismatched statistics in the Access Method Data Set Statistics Block (AMDSB).	Same as for mismatched extents.	No recovery action is required; mismatched statistics do not affect the accessibility of data.

VSAM Object Records (Continued)

*MISCOMPARES Message	Type	Cause	Severity
OTHER	Mismatch of something other than the above fields, e.g. passwords.	Same as for mismatched extents.	Depends on the type of mismatch.

In general, there are two types of data recovery in terms of the currency of the recovered data; repair and reset.

The repair type of data recovery operation restores addressability and access to the most current version of the data. Repair operations are generally used to correct problems such as read and write errors associated with the data itself or with the data description, for example, by assigning alternate tracks.

The reset type of data recovery operation restores addressability and access to a version of the data other than the most recent. Reset operations are generally used to correct logical problems such as a programming error or faulty transactions. Reset is generally the most common form of recovery, probably because of the types of problems encountered and the level of data available for recovery. An example of a reset operation is the dump (copy) and restore of a volume.

The following list of utility programs, whether Access Method Services, VSAM, or system, shows the type(s) of data recovery (or analysis) each program can undertake:

EXPORT/IMPORT	Reset
REPRO	Reset/repair
EXPORTRA/IMPORTRA	Repair (recoverable catalogs only)
COPY and RESTORE	Reset
LISTCRA (COMPARE)	Analysis (recoverable catalogs only)
DELETE UCAT(FORCE)	Reset/repair
ALTER REMOVEVOLUMES	Reset/repair
DELETE SPACE(FORCE)	
NOSCRATCH	Reset/repair
DELETE CLUSTER	
NOSCRATCH	Reset/repair
VTOC UTILITY(SUPERZAP)	Reset/repair
DELETE SPACE (FORCE)	Reset/repair
VERIFY	Repair
RESETCAT	Reset/repair (recoverable catalogs only)

REPRO can be used to create a backup copy of the catalog. The catalog can be reloaded (using REPRO) from this backup copy to minimize the amount of work required to accomplish recovery. The usefulness of a backup copy depends on what modifications have been made to the catalog since the backup was created. Subsequent to taking a backup copy will affect the usefulness of the backup catalog:

Altering the amount of space controlled by the catalog. The volume entry in the backup catalog is no longer valid and will mismatch with the catalog recovery area.

Defining, altering, or deleting data sets. The volume entry and some of the data set entries in the backup copy are no longer valid.

Suballocating space to a VSAM data set. The volume space map is invalidated which, in itself, is not serious. However, the data set entry is also invalidated, a serious problem. (See "Ensuring Accessibility of Secondary Extents" for more information.)

Several of the recovery procedures use volume restore. If this is indicated, one or the other of the following must be true:

- The volume being restored does not contain multivolume data sets.
- If a volume does contain a portion of a multivolume data set, all volumes which contain portions of those multivolume data sets are treated as a single unit; that is, if a volume restore is required, the entire set is restored.

CREATING AND CATALOGING OBJECTS

VSAM uses catalogs as a central information point for all VSAM data sets and the direct-access volumes on which they are stored. You can define a VSAM object in a VSAM catalog only by using the Access Method Services DEFINE command. Additionally, you can use the DEFINE command to define nonVSAM objects in a VSAM catalog.

When you issue the DEFINE command to catalog an object, Access Method Services builds a catalog entry to describe the object. The VSAM objects you can define are:

- User catalog, which is a collection of information about nonVSAM objects, and about VSAM objects that reside on the volumes owned by the user catalog. You can create a user catalog at any time. A pointer to the user catalog is put in the system's master catalog. A master catalog is created during the SYSGEN process. The VSAM master catalog is the OS/VS2 system's primary catalog.
- Data space, which is direct-access device space used for a catalog, for VSAM clusters, for page spaces, and for alternate indexes. VSAM controls the allocation of space within each data space.
- Cluster, or VSAM data set, which is a collection of user-data records. There are three types of cluster data organization:
 - Entry-sequenced, or sequential, in which data records are read or written sequentially from one end to the other (first to last for writing, either direction for reading).
 - Key-sequenced, or indexed, in which a data record is read or written based on its key value. A key is a field, shorter than and within the record, that identifies the record.
 - Relative-record, or direct, in which a data record is read or written based on its relative-record number—its displacement, in records, from the beginning of the cluster.
- Alternate index, which allows you to read and write data records in an entry-sequenced or key-sequenced cluster (called the base cluster) based on an alternate key.
- Path, which is a data-set name for the combination of an alternate index and its base cluster, or an alias for a VSAM cluster. A path entry can be password protected.
- Page space, which is an amount of direct-access device space to be used exclusively by the OS/VS2 system.

The nonVSAM objects you can define are:

- NonVSAM data set, which is a collection of data records with sequential, partitioned, direct, or indexed-sequential data organization (that is, not VSAM data organization).
- Generation data group, which is a collection of nonVSAM data sets that are grouped together in a time-dependent manner.
- Alias, which is an alternate name for a user catalog or nonVSAM data set.

Before an object can be defined, you must have a VSAM catalog in which to define the object. Every OS/VS2 system must have a VSAM master catalog before initial program load (IPL) is done for the system.

When you define an object, you specify attributes to be associated with it. The attributes include, for example, any passwords required to use data and how space is to be allocated. After the object is defined, it can be processed with other Access Method Services commands and with the user's VSAM program. After a cluster is defined, for example, you can load data records into it by using the REPRO command.

VSAM clusters, alternate indexes, page spaces, and catalogs are stored in data spaces. You can allow a data space to contain many VSAM objects by simply not specifying otherwise (the usual case and the default), or you can restrict a data space so that it contains only one VSAM object. In the usual case, you define a data space first, then define the data sets. To define a VSAM object as the only one in its data space, specify the parameter UNIQUE when you define the data set. When the object is defined, VSAM allocates a unique data space and assigns it to the object.

Refer to the chapter: "Command Format" for a description of the command format for DEFINE USERCATALOG. Examples of defining user catalogs are also included.

PREVENTING DUPLICATE NAMES IN CATALOGS

With multiple catalogs (the master catalog, user catalogs, and OS/VS CVOLs), you should take care so that a data-set name in one catalog is not duplicated in another catalog. It is possible to have the same data-set name in more than one catalog.

Access Method Services prevents you from cataloging two objects with the same name in the same catalog, and from altering the name of an object so that its new name duplicates the name of another object in the same catalog. However, this does not prevent duplication of names from one catalog to another. A following section, "Order of Catalog Use: DEFINE," describes the order in which one of the catalogs available to the system is selected to contain the to-be-defined catalog entry. When you define an object, you should ensure that the catalog the system selects is the catalog you want the object's entry in.

Additionally, duplication is not prevented when a user catalog is imported into a system; no check is made to determine whether the imported catalog contains a name that another catalog already in the system contains.

JCL AND DYNAMIC ALLOCATION (DEFINE)

Defining an object does not normally require that a volume be mounted, because the availability of space in data spaces can be determined by examining the volume information in the catalog. A volume must be mounted whenever its VTOC has to be consulted or modified, that is, when a data space, a unique data set, or a catalog is being defined. If the catalog in which the object is being defined is recoverable, the volume containing the catalog recovery area to be updated must be mounted. (Refer to the section "Restoring Catalog Entries After System Failure" for a description of the catalog recovery volume.) Before a system component, such as VSAM or Access Method Services, can process a data set or volume, the OS/VS system allocates the data set or volume to the system component. Usually, the data set or volume is allocated to the system component when the OS/VS Scheduler processes your job control language (JCL) DD statements.

However, when a data-set name or volume-serial number is specified with an Access Method Services command, but its data set or volume has not been described with a DD statement, the data set or volume is allocated dynamically. The information

describing a data set that the DD statement would have provided is extracted from the data set's entry in the catalog. This type of allocation is called dynamic allocation. The data-set's volume should be mounted as permanently resident or reserved to ensure successful dynamic allocation.

If JCL is used to allocate a volume, include a DD statement of the form:

```
//ddname DD UNIT=(type[,unitcount][,DEFER]),  
// VOLUME=SER=(serial[,serial2,...]),DISP=OLD
```

- UNIT= can specify a device type and volume-serial number for any direct-access storage device that is supported by your OS/VS system.
- VOLUME=SER= indicates the serial number(s) of the volume(s) on which data spaces are to be allocated. When more than one volume is indicated, all volumes must be of the same device type.
- DISP=OLD is required because VSAM allocates space to VSAM objects (that is, catalogs, clusters, alternate indexes, and page spaces). When you code DISP=NEW, the OS/VS system allocates space to the object. The space that the system allocates is not recorded in the object's catalog entry. Unless the object is a nonVSAM data set, the space allocated by the system is not available for the object's use.

ORDER OF CATALOG USE: DEFINE

The order in which catalogs are used when an entry is defined is:

- If a catalog is specified in the CATALOG parameter, that catalog is selected to contain the to-be-defined entry.
- When a nonVSAM generation data group (GDG) data set is defined, the catalog containing the GDG base is selected to contain the to-be-defined nonVSAM entry. Otherwise,
- The first user catalog specified in the current job step (STEPCAT) or, if none is specified for the job step, the first user catalog in the current job (JOB CAT) is selected to contain the to-be-defined entry. (The section "JOB CAT and STEPCAT DD Statements" describes how you can specify a user catalog for the current job step or current job.) Otherwise,
- If no user catalog is specified for the current job step or job, and the entry's name is a qualified name, and the first qualifier (that is, the first one to eight characters before the period) is the same as:
 - the name of a user catalog, or
 - the alias of a user catalog,the catalog so identified is selected to contain the to-be-defined entry.
- If no catalog has been identified, either explicitly or implicitly, VSAM defines an object in the master catalog.

THE DEFINE COMMAND: HOW PARAMETERS ARE ORGANIZED

When you define a catalog, cluster, or alternate index, you can specify attributes in several different ways. The parameter set for DEFINE USERCATALOG, CLUSTER, and ALTERNATEINDEX is directly related to the way in which the attributes are stored in the VSAM catalog. The catalog entries which describe a VSAM catalog, cluster, or alternate index are:

- The cluster entry, which describes attributes pertaining to the cluster or catalog as a whole; these are primarily protection attributes.
- The alternate index entry, which describes attributes pertaining to the alternate index as a whole; these are primarily protection attributes.
- The data entry, which describes attributes pertaining to the data component of a catalog, cluster, or alternate index.
- The index entry, which describes attributes pertaining to the index component of a catalog, key-sequenced cluster, or alternate index.

When you specify attributes as parameters of USERCATALOG, CLUSTER, or ALTERNATEINDEX, they are defined in the VSAM catalog in the cluster or alternate index entry if they pertain to that entry; for example, protection attributes. However, when you specify attributes as parameters of USERCATALOG, CLUSTER, or ALTERNATEINDEX which pertain to the data and/or index entries, Access Method Services directs the attributes to the proper entry. For example, if you specify the WRITECHECK attribute as a parameter of CLUSTER, it will not be defined in the CLUSTER entry but will be entered in the DATA entry and the INDEX entry in the catalog. When you specify attributes as parameters of DATA or INDEX, they are defined in the VSAM catalog directly in the data or index entry. You may use the LISTCAT command with the ALL option to list catalog entries and determine where the various attributes are stored in the catalog. "Appendix B: Interpreting LISTCAT Output Listings" describes the various attributes for each type of entry.

Except for protection attributes, if you specify the same attribute as a subparameter of USERCATALOG, CLUSTER, or ALTERNATEINDEX and also as a subparameter of DATA and/or INDEX, the value of the attribute specified at the DATA/INDEX level overrides the value of the attribute specified at the USERCATALOG, CLUSTER, or ALTERNATEINDEX level.

Since Access Method Services directs the attributes you specify to the proper entry, you do not need to specify any attributes as parameters of DATA or INDEX. However, you may wish to specify attributes as parameters of DATA and INDEX if you wish to:

- Specify a name for the data and index components of a cluster or alternate index.
- Specify unique protection attributes for the data and index components of a cluster or alternate index.
- Specify space allocation for the data component only or for the data and index components separately of a catalog, key-sequenced cluster, or alternate index.
- Specify volumes of a different device type for the data and index components of a key-sequenced cluster or alternate index.
- Specify a different variation of the same attribute for the data and index components of a catalog, key-sequenced cluster, or alternate index. For example, you may wish to specify WRITECHECK for the data component and NOWRITECHECK for the index component or vice versa.

The level at which you specify attributes can also be affected if you use the MODEL parameter of the DEFINE command (see "How to Use One Object as a Model for Another Object").

HOW TO USE ONE OBJECT AS A MODEL FOR ANOTHER OBJECT

You can use the entry of an already-defined VSAM object (that is, an already-defined alternate index, catalog, cluster, or path) as a model for the definition of another object of the same type. When one entry is used as a model for another, its attributes are copied as the new entry is defined.

You can override some of the model's attributes by explicitly specifying them with the DEFINE command. If you do not want to change or add any attributes, you need only supply the entryname of the object being defined and the MODEL parameter. (When you define a catalog, you must also specify the catalog's volume and space information, because the catalog, when defined, owns the volume. When you define an alternate index or a path, you must also specify the object to which it is related.)

- If MODEL is specified as a parameter of USERCATALOG, PAGESPACE, or PATH:
 - the attributes of the model are copied for the object being defined, then
 - any attributes explicitly specified as parameters of your DEFINE command override those of the model.
- If MODEL is specified as a parameter of CLUSTER or ALTERNATEINDEX (that is, at the cluster level) but is not specified as a subparameter of the DATA or INDEX parameter:
 - the attributes of the model are copied for the cluster or alternate index entry of the object being defined, then
 - any attributes explicitly specified as parameters of CLUSTER or ALTERNATEINDEX override those of the model for the cluster or alternate index entry.
 - the attributes of the model's data and index components are copied for the object's data and index components, then
 - attributes explicitly specified as parameters of CLUSTER or ALTERNATEINDEX are propagated to the data and index components, overriding those of the model, then
 - attributes explicitly specified for the object's data and index components (that is, specified as subparameters of the DATA or INDEX parameter) override the previous two steps.
- If MODEL is specified as a subparameter of DATA or INDEX as well as a parameter of CLUSTER or ALTERNATEINDEX:
 - the attributes of the CLUSTER or ALTERNATEINDEX model are copied for the cluster or alternate index entry of the object being defined, then
 - any attributes explicitly specified as parameters of CLUSTER or ALTERNATEINDEX and override those of the model for the cluster or alternate index entry.
 - attributes explicitly specified as parameters of CLUSTER or ALTERNATEINDEX are propagated to the object's data and index components, then
 - attributes of the model specified for the data or index component are copied (that is, the model specified with the MODEL subparameter of the DATA or INDEX parameter), overriding the previous step, then

- attributes explicitly specified for the data and index components are copied (that is, the attributes specified with subparameters of the DATA or INDEX parameter), overriding the previous two steps.

The MODEL parameter is designed to let you easily define data sets that are essentially identical. By explicitly specifying certain parameters (for example, protection attributes), you can override any attributes of the model you wish to change. When you use the MODEL parameter, you should take care to ensure that your job is not terminated because of allocation problems when you explicitly do any of the following:

- Specify a different type of device with the VOLUMES parameter.
- Change the length or position of the keys with the KEYS parameter.
- Change the size of records, buffer space, or control intervals with the RECORDSIZE, BUFFERSPACE, or CONTROLINTERVALSIZE parameters.
- Change the type of cluster (that is, entry-sequenced, key-sequenced, or relative-record), the type of alternate index (that is, key-pointer or RBA-pointer), or the allocation-type of the object (that is, unique or nonunique).
- Change the unit of allocation with the CYLINDERS, RECORDS, or TRACKS parameters.

When you explicitly specify any of the above parameters for your to-be-defined object, you might have to make corresponding changes to other related parameters.

DEFINING A CATALOG

The DEFINE command can be used to define user catalogs.

A master catalog, which is the system catalog, is established at system-generation time. The volume on which the master catalog is defined must be permanently mounted. The master catalog is selected for use at initial program load (IPL) time. Because the master catalog is created at system-generation time and because only one master catalog is allowed on a system, you cannot issue DEFINE MASTERCATALOG to create a master catalog in a VS2 system. If you issue DEFINE MASTERCATALOG, VSAM creates a user catalog. You can, however, establish a user catalog as a master catalog at IPL time. See "Appendix H: Changing the Attributes of the VSAM Master Catalog" for a description of this procedure.

In a system with the Mass Storage System, a user catalog can be defined on a mass storage volume. If it is, it is staged and bound when it is opened (that is, it is retained until closed on the direct-access storage staging drive to which it is staged). A catalog can contain entries for both objects stored on mass storage volumes and objects stored on direct-access storage volumes.

User catalogs are pointed to by the master catalog. The master catalog contains a user-catalog entry that identifies the user catalog's volume. Each user catalog resides on a different volume. All VSAM objects on a volume must be defined in the volume's user catalog. A user catalog can control (that is, catalog the objects of) more than one volume.

When you define a user catalog, you can specify passwords for it even when the master catalog is not password-protected. To prevent an unauthorized user from deleting your user catalog, specify a password for it (the catalog's master password is

required to delete it, but any password specified for the catalog results in a master password). (See the MASTERPW, CONTROLPW, UPDATEPW, and READPW descriptions.)

You can delete a user catalog and redefine it in the same step, but you cannot use it in that step. A TSO session is considered one step so you would have to log off and log on to use a catalog that you had deleted and redefined.

Note: When a new user catalog has been defined, the first object cataloged causes the initialization of the entire catalog to take place. It also reserves the volume on which the catalog resides.

HOW SPACE IS ASSIGNED TO A CATALOG

The structure of a VSAM catalog is described in Planning for Enhanced VSAM Under OS/VS. Familiarity with this structure will be of help to you in understanding the discussion that follows.

When a catalog is defined, the catalog is the first VSAM object contained on the volume. VSAM first builds a data space on the volume and then builds the catalog itself (see the section "Defining a Data Space" for further information about VSAM data spaces). You must specify the size of the data space by using the CYLINDERS, RECORDS, or TRACKS parameter as a subparameter of USERCATALOG. For example,

```
DEFINE USERCATALOG -  
  TRACKS(prim secn)...)
```

where:

- TRACKS specifies the number of tracks to be allocated to the data space.

You may optionally specify the size of the catalog itself by using the CYLINDERS, RECORDS, or TRACKS parameter as a subparameter of DATA, or DATA and INDEX. For example,

```
DEFINE USERCATALOG -  
  (TRACKS(prim secn)... ) -  
  DATA -  
  (RECORDS(prec srec)...)
```

where:

- TRACKS specifies the number of tracks to be allocated to the data space.
- RECORDS specifies the number of records to be contained in the low-key range of the catalog's data component (see the section "Estimating the Catalog's Space Requirements" for a discussion of this value).

The data space is built for the catalog's exclusive use when the DATA and INDEX parameters do not include space subparameters. If the DATA or DATA and INDEX parameters do include space subparameters, part of the data space can contain suballocated VSAM data sets.

To share the data space with other VSAM data sets, you can specify space for the data and, optionally, index components of the catalog. The space specified as a subparameter of DATA and INDEX determines the amount of space used by the catalog; the space specified as a subparameter of USERCATALOG determines the amount of space in the data space. If space is specified for the low-key range of the data component (using the DATA parameter) but not for the index component, VSAM determines the amount of space required for the catalog's data component high-key range, the sequence set, and for the index component.

The following is a summary of the three ways in which space can be specified when defining a catalog.

- If DATA space subparameters are not coded, the entire data space is allocated to the catalog in multiples of allocation units. Any remaining space that is not a multiple of an allocation unit is available for suballocation by other VSAM data sets. (INDEX space subparameters are invalid unless DATA space subparameters are also specified.)
- If both DATA and INDEX space subparameters are coded, the sum of the space subparameters determines the amount of space available to contain the catalog. VSAM determines the proportion of this amount of space to be allocated to the catalog's data and index components. The remainder of the data space is available to contain suballocated VSAM data sets.
- If DATA space subparameters are coded but INDEX space subparameters are not coded, the amount specified determines the size of the low-key range of the catalog's data component. VSAM then determines the amount of space required for the high-key range of the data component, the sequence set, and the high-level index and adds it to the DATA space subparameter specification. That is, VSAM determines the additional space needed based on the amount of space specified for the low-key range of the catalog's data component. The remainder of the data space (that is, the space that is not allocated to either the catalog's data or index components) is available to contain suballocated VSAM data sets.

For example, if you specified, as a subparameter of USERCATALOG, 100 cylinders of a 3330 Disk Storage Device (that is, the primary allocation for the catalog's data space is 100 cylinders) and for the catalog's data and index components 5 cylinders (that is, the sum of the DATA and INDEX primary allocation amounts is 5 cylinders), 95 cylinders remain in the data space and are available for suballocated VSAM data sets.

Catalog space is allocated in tracks (even when you specify cylinders or records). The number of tracks is a multiple of an allocation unit. An allocation unit is the size of a control area, with one track for the control area's replicated sequence-set record. The size of an allocation unit is:

- 3 tracks for a 3380, 3375, 3350, 3330, 3330 model 11, or 2305 model 2 (2 tracks for the control area and 1 track for the sequence set record)
- 4 tracks for a 2305 model 1 (3 tracks for the control area and 1 track for the sequence set record).
- 5 tracks for 2314/2319 or 3340/3344 (4 tracks for the control area plus 1 track for the sequence set record).

When the amount of space you specify is not a multiple of the device's allocation unit, the amount of space is rounded downward to a multiple and the additional tracks are not used. If you specify less than the minimum allocation unit, the amount of space is rounded upward and a minimum-sized catalog is allocated.

One allocation unit is always allocated to the index set of the index. Ten percent of the remaining allocation units (or at least one allocation unit) is allocated to the high-key range of the data component. The rest of the allocation units are allocated to the low-key range. The minimum catalog size is three allocation units: one each for the index set, the high-key range, and the low-key range.

For example, assume that one cylinder on a 3330 is specified for the entire catalog (1 cylinder = 19 tracks):

- The number of tracks is rounded down to 18 tracks, which is six allocation units of three tracks each (the nineteenth track is not used and not available unless the catalog's data space is suballocated).
- One allocation unit (three tracks) is allocated to the catalog's index set.
- Because 10% of the remaining 15 tracks is less than three tracks, one allocation unit is allocated to the high-key range of the data component. The first track contains a replicated sequence-set record. The next two tracks contain a control area with 40 control intervals, which can hold a maximum of 400 records. (Records in the high-key range are 47 bytes long.)
- The remaining four allocation units (12 tracks) are allocated to the low-key range of the data component. Each allocation unit includes one track for the replicated sequence-set record and two tracks for a control area with 40 control intervals, which can hold 40 records. (Records in the low-key range are 505 bytes long.) The low-key range can hold a maximum of 160 records. From 12 to 15 of these records (13 for a 3330) are used to describe the catalog itself.

When the maximum number of records is reached in the low-key or high-key range, the catalog must be extended to hold additional records. When you define the catalog, you can specify an amount of space for secondary allocation.

It may be desirable in some cases to have the high-key range allocated 20% of the data component space, rather than the normal 10%. If the SPACE subparameter is specified on both DATA and INDEX, and the specified INDEX quantity is exactly 20% of the specified DATA quantity, then the high-key range will be allocated 20% of the space available for the data component.

```
DEFINE USERCATALOG( -  
  NAME(USERCAT) -  
  FILE(DDX) -  
  VOLUME(USER01) - -  
  CYLINDERS(60 5)) -  
  DATA( -  
  CYLINDERS(50 5)) -  
  INDEX( -  
  CYLINDERS(10))
```

Sixty cylinders are to be allocated to the catalog's data space. Since space is specified on both DATA and INDEX, the index value is exactly 20% of the data value, the high-key range will be allocated 20% of the space available for the data component.

IDENTIFYING THE CATALOG'S VOLUME

When you define a catalog, you can use a JCL DD statement to identify and mount the volume on which the catalog is to reside. Instead of using a JCL DD statement, you can rely on dynamic allocation to identify the catalog's volume. The volume must be mounted as permanently resident or reserved. If JCL is used, include a DD statement of the form:

```
//ddname DD DISP=OLD,UNIT=(type[,DEFER]),  
// VOLUME=SER=serial
```

The volume on which the catalog and its data space are to be allocated must be mounted, because the data space must be described in a data-set control block (DSCB) that is stored in the volume's VTOC (volume table of contents). The DSCB is added to the VTOC and the volume-ownership information, also in the

VTOC, is reset to indicate that the volume is owned by a VSAM catalog. VSAM determines whether the space is available on the volume and, if so, formats a DSCB for the data space.

NONVSAM DATA SETS AND THE CATALOG'S UPDATE PASSWORD

When you expect the user catalog to contain nonVSAM data sets, you should take care when you specify passwords for the catalog. When you specify an update password for the catalog, Access Method Services requires that you supply the password whenever you (or another user) want to alter or delete one of the catalog's nonVSAM data sets. You can protect the catalog against unauthorized usage by specifying a master password. See the UPDATEPW description under DEFINE USERCATALOG for more details about how a catalog's update password is used.

ESTIMATING THE CATALOG'S SPACE REQUIREMENTS

Before you define a catalog, you should estimate the amount of space needed for the catalog's data component. To determine the approximate number of records required for the low-key range of the catalog, use the worksheet in Figure 3:

Variable Quantities	Formulas	Estimates
Basic requirement = 10 records		10
A = number of key-sequenced clusters	Ax3	
A ¹ = number of key-sequenced clusters with alternate indexes to be upgraded	A ¹ x4	
B = number of entry-sequenced clusters	Bx2	
B ¹ = number of entry-sequenced clusters with alternate indexes to be upgraded	B ¹ x3	
C = number of relative-record clusters	Cx2	
D = number of alternate indexes	Dx3	
E = number of path entries	E	
F = number of nonVSAM data set entries	F	
G = number of generation data group entries	G	
H = number of alias entries	H	
I = number of pagespaces	I*2	
J = number of volumes, depending on device type, owned by the catalog:		
J ¹ = number of 2305 volumes	J ¹ x2	
J ² = number of 2314/2319 volumes	J ² x3	
J ³ = number of 3330 and 3340/3344 volumes	J ³ x4	
J ⁴ = number of 3330 Model 11 and 3350 volumes	J ⁴ x6	
K = for each key-sequenced cluster and alternate index (KSDS) with space on more than two volumes, add "1" for each additional group of one to five volumes:		
K ¹ = number of KSDSs with 3 to 7 volumes	K ¹	
K ² = number of KSDSs with 8 to 12 volumes	K ² x2	
K ³ = number of KSDSs with 13 to 17 volumes	K ² x3	
L = for each entry-sequenced cluster and relative-record cluster (ESDS) with space on more than five volumes, add "1" for each additional group of one to eight volumes:		
L ¹ = number of ESDSs with 6 to 13 volumes	L ¹	
L ² = number of ESDSs with 14 to 21 volumes	L ² x2	
M = for each group of four data spaces on a volume, add "1"	M	
N = number of entry records required for the catalog's data component (total of A through M above)	N	

Figure 3. Worksheet for Estimating a Catalog's Space Requirements

When you define the catalog's space parameters, you may use one of the formats shown below:

- If you wish to assign all of the space to the catalog itself, you may specify:

```
DEFINE USERCATALOG -
  (TRACKS{prim secn}...)
```

where:

- prim = the amount of space for the primary extent (that is, allocation) of the catalog's data space. You can calculate the minimum value of prim as:

$$\text{prim} = (aN + b) \text{ tracks}$$

where N is the value derived from the worksheet (Figure 3) and the values of a and b are derived from the table below:

Catalog resides on:	a	b
2305 Model 1	0.1	5
2305 Model 2	0.09	3
2314 or 2319	0.1364	5
3330 or 3330-11	0.09	3
3340 or 3344	0.125	5
3350	0.0667	3

The calculated value of prim should be rounded upward to a whole number of tracks.

- `secn` = the amount of space for each secondary extent of the catalog's data space.

This example shows the space specification in terms of TRACKS using the worksheet in Figure 3. You may also specify the value for `prim` and `secn` as a subparameter of CYLINDERS.

2. If you wish to have other VSAM objects in the catalog's data space, you may specify:

```
DEFINE USERCATALOG -  
  (TRACKS(prim secn)... ) -  
  DATA -  
  (RECORDS(prec srec ... )
```

where:

- `prec` = the number of catalog entry records in the low-key range, or the value of N from the worksheet (Figure 3). VSAM uses this value to calculate the total primary allocation requirements for the catalog itself.
- `srec` = the number of catalog-entry records to be used by VSAM to calculate the catalog's secondary-extent allocation.
- `prim` = the total amount of space for the primary extent of the catalog's data space. The amount of space assigned to the catalog itself is taken from the `prim` space as calculated by VSAM from the value specified by `prec`. You can determine the amount of space that VSAM calculates for the catalog itself using the following algorithm:

$(aN + b)$ tracks

as described for format-1 above. When you specify a value for `prim` larger than the requirements for the catalog itself, the remainder of the space is available for the suballocation of other VSAM objects. Using this format, you may also specify the value for `prim` as a subparameter of CYLINDERS rather than TRACKS.

- `secn` = the amount of space to be used for each secondary extent of the catalog's data space.

The Catalog's Secondary Allocation Amount

You should specify a secondary allocation amount for the catalog's data component. You cannot use the ALTER command to add or change a secondary allocation amount once the catalog is defined.

The secondary allocation quantity is allocated each time one of the catalog's key ranges needs more space—up to a total of thirteen times. The secondary allocation quantity you specify

for the catalog's data component applies to the catalog's high-key range and low-key range (including their sequence set tracks). When the high-key range or low-key range runs out of space, the specified secondary allocation quantity is allocated to the key range that requires the space (when both key ranges run out of space at the same time, two separate extents are allocated: one for each key range, and each equal to the secondary allocation quantity).

You use the secondary allocation to minimize the size of the low-key range and to provide for an extension of the high-key range: after the catalog's index set is allocated, 10% of the remaining space (or one allocation unit, whichever is larger) is allocated to the high-key range; the remainder of the catalog's space is allocated to the low-key range. However, the high-key range might achieve only 25% utilization. Rather than specify a primary allocation quantity large enough to allow for low utilization of the high-key range and, therefore, create an unnecessarily large low-key range, you can specify a secondary allocation. The secondary allocation quantity is allocated to the high-key range when it requires space, without affecting the size of the low-key range. (The same applies to the low-key range, when it requires space.)

The high-level index of a catalog never requires extension.

Space Allocation for the Catalog Recovery Area (CRA)

Whether or not a catalog is defined with the RECOVERABLE option has a bearing on how space is assigned in the catalog. For a recoverable catalog, VSAM creates a catalog recovery area (CRA) on each volume owned by the catalog. The primary allocation for a CRA is the number of tracks equivalent to one cylinder. VSAM ensures that these tracks are contiguous, but they may or may not start on a cylinder boundary. On the volume on which the catalog resides, the primary allocation for the CRA is allocated from the catalog's data space.

On volumes without a VSAM catalog, a primary space allocation is obtained by VSAM for the CRA when the first data space is defined on the volume. If the first data space is defined via the UNIQUE attribute of DEFINE CLUSTER or ALTERNATEINDEX, VSAM obtains one cylinder of space on the volume for the CRA. In other words, the CRA's primary and secondary allocation is added to the space specified via space parameters (CYLINDERS, RECORDS, or TRACKS) of your DEFINE command.

If the first data space is defined via DEFINE SPACE, VSAM first attempts to obtain space for the CRA's primary and secondary allocation as before. However, if VSAM fails to obtain the space, the CRA will be allocated from the space specified via the space parameters. If the first DEFINE SPACE for a volume specifies the CANDIDATE attribute, VSAM obtains one cylinder for the CRA's primary and secondary allocation.

Once a primary allocation of the CRA is filled, it can expand to include a maximum of 15 additional extents. The allocation for each secondary extent is one cylinder.

DEFINING A DATA SPACE

The DEFINE command can be used to define VSAM data spaces or to reserve volumes for future VSAM use.

Refer to the chapter: "Command Format" for a description of the command format for DEFINE SPACE. An example of defining a VSAM data space is also included.

VSAM DATA SPACES ON A VOLUME

A VSAM data space is space on a direct access volume that is owned and managed by VSAM (as opposed to space managed by OS/VS using Direct-Access Device Space Management—DADSM). All of the data spaces (and VSAM data sets) on a volume are defined in the same VSAM catalog. When you define the volume's first data space (or you specify the volume as a candidate to contain VSAM objects) you are, in effect, giving control over the volume to the catalog that contains the data space entry. All future data spaces and VSAM objects on that volume must be defined in the defining catalog (that is, the catalog that describes the volume's first-defined data space). Remember that when you define a catalog you cause the first data space to be defined on the catalog volume.

A VSAM data space may include several extents on the same volume, but it cannot span volumes. The volume that contains the data space is owned by the catalog that contains the catalog entry for the space. A VSAM data space can take up all or part of the space of a direct-access volume. It cannot take up space on more than one volume. There can be more than one data space on a volume. In a system with the Mass Storage System, a data space can be defined on a mass storage volume.

When you define a data space, you must specify the volume that is to contain the data space. When you specify more than one volume, a data space of the size you specify is allocated on each volume; each volume comes under the control of the defining catalog. Access Method Services creates a volume entry in the catalog to describe each volume on which one or more data spaces have been defined.

If there is space available and there is no volume ownership conflict, a VSAM data space is allocated on the volume. When the data space is the first data space on a volume, VSAM ownership is indicated in the VTOC. The volume's VTOC is modified so that an ownership indicator is set in the format-4 DSCB. A format-1 (identifier) DSCB describing the data space is also written into the VTOC. When the data space is the first data space on a volume owned by a recoverable catalog, VSAM also allocates a recovery space on the volume. Refer to "Space Allocation for the Catalog Recovery Area (CRA)" for more information on CRA space allocation.

Allocating the Volume

You can use a DD statement to allocate the volume on which the data space is to be defined. If you do not use a DD statement to allocate the volume, the volume identified with the VOLUMES parameter is dynamically allocated. The volume must be mounted as permanently resident or reserved.

When you use a DD statement to allocate the volume on which the data space is to be defined, it should be in the form:

```
//ddname DD DISP=OLD,  
//          UNIT=(devtype[,unitcount][,DEFER]),  
//          VOLUME=SER=(volser1[,volser2,...])
```

VSAM OBJECTS IN A DATA SPACE

All VSAM objects (that is, catalogs, clusters, alternate indexes, and page spaces) are stored in VSAM data spaces. A unique component is defined as the only component in its data space. VSAM allocates the data space for a unique component when you define the data set. A nonunique object can share a data space with other VSAM objects. That is, a catalog, clusters, page spaces, and alternate indexes might also be in the data space. To define a data set that can share a data space with other data sets, a data space must have been defined beforehand in which to suballocate space for the data set.

A data space can contain many nonunique data sets, and a data set can be stored in more than one data space, on the same volume or on different volumes.

SPACE ASSIGNMENT TO VSAM OBJECTS

When VSAM allocates space for a new data set or extends the space for an existing data set, it allocates the space from that available in a data space. For a new data set, if not enough space is available, the DEFINE fails, and you must define a data space with sufficient space.

For an existing data set, if not enough space is available, VSAM tries to extend the data space. If the data space cannot be extended (it may already have the maximum number of extents or there may not be enough space on the volume), VSAM tries to set up a new data space on the same volume as the old.

If the data space cannot be extended because a suballocated data set's secondary allocation amount is greater than the data space's secondary allocation amount, VSAM builds a new data space: its primary and secondary allocation amounts are the same as the data set's secondary allocation amount. Consequently, the total amount of space allocated to VSAM data spaces on a volume might exceed the predicted limit. If there is not enough space on the volume and you identified another volume for expansion when you defined the data set, VSAM tries to set up a data space on that volume.

DEFINING A CLUSTER

The DEFINE command can be used to define a catalog entry and allocate space for an indexed, nonindexed, or relative-record cluster. When a DEFINE command is used to define a cluster, as many as three entries are created in the catalog: an entry for the cluster, its data component, and, if the cluster is indexed, its index component.

Attributes of components can be specified separately from the attributes of the cluster. If attributes are specified for the cluster as a whole and are not specified for the components, the attributes of the cluster (except for its passwords and other protection attributes) apply to its components. If an attribute that is applicable to the data or index component is specified for both the cluster and the component, the explicit component specification overrides that specified for the cluster.

Note: A qualified data set name and an unqualified data set name cannot be placed in the same catalog if the first qualifier of the qualified name is the same as the unqualified name. For example, PAYROLL and PAYROLL.DATA cannot exist on the same catalog.

You specify a name for the cluster when you define it and generally give this name as the dsname in JCL. You can optionally name the components of a cluster. You may wish to name the data component of an entry-sequenced cluster or the data and index components of a key-sequenced cluster in order to make it easier to process them individually. For instance, you may open the index of a key-sequenced data set and process it as data. This processing is described in QS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications. If you wish to use the ALTER command to modify attributes of any component, you must identify that component by name. If you do not specify a name for the components of a cluster, VSAM generates a name in the form:

VSAMDSET.Tbbbbbbb.DFDyyddd.Taaaaaaa.Tbbbbbbb

where:

- yyddd is the date (year and Julian day)
- aaaaaaabbbbbbb is the timestamp value

Refer to the chapter "Command Format" for a description of the command format for DEFINE CLUSTER. Examples of defining clusters are also included.

SPECIFYING CLUSTER INFORMATION

You specify all of the necessary descriptive information and the performance, security, and integrity options when you create a cluster. Some of the information you specify can apply to either the data or the index or to both. You can specify information for the cluster as a whole in the CLUSTER parameter. Information for data only or for index only goes with the corresponding parameter (DATA or INDEX).

Descriptive Information

Descriptive information includes:

- Type of data organization
- Whether the cluster is reusable (that is, for temporary storage of data)
- Average and maximum lengths of data records
- Length and position of the key field in the records of a key-sequenced data set
- Identity of the catalog in which to define the cluster
- Identity of the volume(s) on which to allocate space for the cluster
- Amount of space to allocate for the cluster and whether the cluster is to reside in a separate data space
- Least amount of I/O-buffer space a processing program will provide to process the data set

With a multivolume key-sequenced data set, you may assign data to the various volumes according to ranges of key values. For example: if you have three volumes you might assign records with keys A-E to the first volume, F-M to the second, and N-Z to the third. Key-range allocation facilitates processing the data set with only one of the volumes mounted: you know which volume contains what records. All of the volumes specified for the cluster's data component or for the index component must be of the same device type. However, the data component's device type can be different than the index component's device type.

If a component is unique (that is, is to reside alone in a separate data space), the data space is not defined before the cluster is defined. The data space is created as the cluster is defined; in this case, the volume(s) must be mounted.

If you don't specify the space to be provided for buffers, VSAM determines control interval size first and then sets buffer-space amount equal to the size of two data control intervals plus, for a key-sequenced data set, one index control interval. If you do specify it, data and index control intervals are limited to sizes such that the buffer space can hold two data control intervals and one index control interval.

If the values you specify for record length and key length require control intervals too large for the buffer space you specify, the DEFINE command will fail.

The relationship between control interval size and least amount of I/O-buffer space is further discussed in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide. If you specify track allocation or specify record allocation for less than one cylinder, space is actually allocated by tracks.

Performance Options Information

Information for performance options (mostly for an indexed cluster) include:

- Whether records can span control intervals
- Whether to replicate index records
- Whether to place the sequence set of an index adjacent to data
- Whether to place the cluster's index on a separate volume from data
- The amount of free space to remain in the data component's control intervals and control areas when the data records are loaded (applies only to a key-sequenced cluster's data component)
- Control interval size for VSAM to use instead of calculating the size itself

These options are discussed in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.

In a system with the Mass Storage System, you can also indicate how a cluster or component that is stored on a mass-storage volume is to be staged.

Protection and Integrity Information

Information for protection and integrity options include:

- Passwords and related information
- Identity of your own authorization routine to verify that a requester has the right to gain access to data
- Identity of an I/O error-handling routine (the exception exit routine) that is branched to before the program's SYNAD exit is taken
- Whether to preformat control areas during data record insertion
- Whether to verify that write operations are without error
- Whether and to what extent to share data among systems, jobs, and subtasks
- Whether to erase the information that a data set contains when you delete the data set

These options are discussed in the chapter "Data Security and Protection" and in Planning for Enhanced VSAM Under OS/VS.

In a system with the Mass Storage System, you can also indicate whether a cluster or component that is stored on a mass storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it.

DEFINING AN ALTERNATE INDEX

An alternate index can be defined over a key-sequenced cluster or over an entry-sequenced cluster. An alternate index cannot be defined to support a reusable cluster, a relative-record cluster, a catalog, another alternate index, or a nonVSAM data set.

Before you can build an alternate index, you must define and load its base cluster (see the DEFINE CLUSTER command), then define the alternate index and relate it to the base cluster. The base cluster and alternate index are described by entries in the same catalog.

You can issue the DEFINE ALTERNATEINDEX command to define an alternate-index entry and allocate space for an alternate index. VSAM uses more than one catalog entry to describe the alternate index (an alternate index is conceptually a key-sequenced cluster):

- An alternate-index entry describes the alternate index as a key-sequenced cluster.
- A data entry describes the alternate index's data component.
- An index entry describes the alternate index's index component.

Attributes of the alternate index's components can be specified separately from the attributes of the alternate index. When attributes are specified for the alternate index as a whole and are not specified for the components, the attributes of the alternate index (except for its passwords and other protection attributes) are propagated to the alternate index's components. The attribute specified for the component overrides that specified for the alternate index.

When you define the alternate index, you identify it with an entryname. The entryname identifies the alternate index when it is described with a JCL DD statement (that is, the entryname is the JCL DD statement's dsname) and when it is processed with another Access Method Services command. You can also name the alternate index's components; a user's program can open and process the alternate index's data or index component as a data set. For example, your program can open the index of an alternate index, then read and write its records. For more details on this kind of processing, see OS/VS Virtual Storage Access Method (VSAM) Options for Advanced Applications.

Refer to the chapter "Command Format" for a description of the command format for DEFINE ALTERNATEINDEX. An example of defining an alternate index is also included.

SPECIFYING ALTERNATE-INDEX INFORMATION

When you define an alternate index, you specify descriptive information and performance, security, and data-integrity options. Some of the information you specify can apply to either the alternate index's data component, or to its index component, or to both—that is, to the alternate index as a whole. You specify information for the alternate index as a whole with the ALTERNATEINDEX parameter and its subparameters. Information for the data component or for the index component is specified with parameters of DATA or INDEX.

Descriptive Information

Information that you specify to describe the alternate index being defined includes:

- The base cluster related to the alternate index.
- Whether the alternate index is reusable (that is, temporary).
- Average and maximum lengths of alternate-index records.
- Length and position of the key field in data records of the alternate index's base cluster.
- Identity of the catalog that is to contain the alternate index's entries (the same catalog that contains the base cluster's entries).
- Identity of the volume(s) on which space is to be allocated for the alternate index.
- Amount of space to allocate for the alternate index and whether the alternate index is to reside in a separate data space.
- The least amount of I/O buffer space that a user's program is to provide when the program processes the alternate index's data.

You can define an alternate index with its records assigned to various volumes according to alternate-key value. Each volume can contain the alternate-index records whose alternate key values are between a low-key value and a high-key value—that is, records whose alternate-key values are within a key range. For example, your alternate index might reside on three volumes. The records with alternate keys A through E are on the first volume, F through M on the second volume, and N through Z on the third volume.

Performance Options Information

Information that you can specify to improve the performance of your alternate index include:

- Control-interval size, instead of allowing VSAM to calculate an optimum size.
- Whether to place the index's sequence set adjacent to the data control areas.
- Whether to place the alternate index's index on a separate volume (or direct-access device) from the alternate-index records.
- The amount of free space to remain in the data component's control intervals and control areas when the alternate-index records are loaded (applies only to the alternate index's data component).

These options are discussed in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.

In a system with the Mass Storage System, you can also indicate how an alternate index or component that is stored on a mass-storage volume is to be staged.

Protection and Integrity Information

Information that you can specify to protect your alternate index's records from other users and from possible system errors include:

- Passwords and related user-verification information.
- Identity of your authorization routine to verify that a user's program is authorized to gain access to the alternate index's records.
- Identity of your I/O error-handling routine (the exception exit routine) that is branched to before the program's SYNAD exit is taken.
- Whether to preformat control areas before the alternate-index records are loaded.
- Whether, and to what extent, to share your alternate index between subtasks, jobs, and systems.
- Whether to erase the information that your alternate index contains when you delete it.

These options are discussed in "Data Security and Protection" and in Planning for Enhanced VSAM Under OS/VS.

In a system with the Mass Storage System, you can also indicate whether an alternate index or component that is stored on a mass-storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it.

DEFINING A PATH

The DEFINE PATH command is used to establish the relationship, the path, between an alternate index and its base cluster. The base cluster and its alternate index must already be defined when you define the path that relates them.

When your program opens a path for processing, both the alternate index and its base cluster are opened. When data in the base cluster is read or written using the path's alternate index, keyed processing is used; RBA processing is not allowed.

You can also use the DEFINE PATH command to establish a password-protected alias for a VSAM cluster. The cluster must already be defined when you define a path as its alias. The path entry, as an alias, allows the cluster another set of protection attributes. You can specify NOUPDATE access for the cluster, which allows you to open the cluster without also opening its upgrade set (that is, all of the cluster's alternate indexes that are to be updated whenever the cluster's data is changed).

Refer to the chapter "Command Format" for a description of the command format for the DEFINE PATH command. An example for defining a path is also included.

DEFINING A NONVSAM DATA SET

The DEFINE command can be used to catalog a nonVSAM data set in a VSAM catalog. Any already existing nonVSAM data set can be introduced into a master or user catalog through the DEFINE command. The result of a DEFINE command when it is used to define a nonVSAM entry is that an entry is created in a master or user catalog; no space is allocated or reserved.

When you define, alter, or delete a nonVSAM data set in a password-protected catalog, the catalog's update- or higher-level password is required.

Refer to the chapter "Command Format" for a description of the command format and an example of the DEFINE NONVSAM command.

DEFINING AN ALTERNATE NAME

The DEFINE command can be used to define an alternate name—an alias—for a nonVSAM data set (even if it is a generation data set). DEFINE ALIAS can also be used to define an alias for a user catalog connector in the master catalog. An alias cannot, however, be defined for a generation data group.

Refer to the chapter "Command Format" for a description of the command format for DEFINE ALIAS. An example is also included.

HOW TO USE AN ALIAS TO IDENTIFY A USER CATALOG

When you define an alias for a user-catalog connector, you should structure the alias so that the catalog's cataloged data sets can be located when the alias is used. If VSAM is searching¹ for a user-specified entryname (that identifies a data set) and does not find its entry in the master catalog or a user catalog identified with the JOBCAT or STEPCAT DD statements, and if the entryname's catalog is not specified, VSAM assumes that:

- The entry resides in a user catalog, and
- The user-catalog's name or alias is the first simple name of the qualified entryname.

If the entryname is ABC.DE.DATA, VSAM searches the master catalog for a user-catalog connector entry (that is, an entry with the name or alias of ABC). If found, VSAM next searches catalog ABC for an entry identified by the name or alias ABC.DE.DATA.

If you identify a catalog with a name of ABC.DE and catalog the entry whose name is ABC.DE.DATA in that catalog, the catalog won't be found when VSAM searches for it. In order for VSAM to find the entryname ABC.DE.DATA (an entry in the catalog whose name is ABC.DE), the catalog's name, ABC.DE, must also be specified (either explicitly in the Access Method Services command that refers to data set ABC.DE.DATA, or as a JOBCAT or STEPCAT DD statement).

You can delete the alias of a user catalog and redefine it in the same step, but you cannot use it in that step. A TSO session is considered one step, so you would have to log off and log on to use an alias of a user catalog that you had deleted and redefined.

¹ For a detailed description of, and the correct search sequence for, how VSAM searches catalogs to locate an entry, see the section "Order of Catalog Use" for the various commands.

DEFINING A GENERATION DATA GROUP

The DEFINE command can be used to create a catalog entry for a generation data group. For further information on generation data groups, see OS/VS2 MVS Data Management Services Guide.

Once a catalog entry for a generation data group has been created, existing data sets can be attached to it as generation data sets. This is accomplished through JCL, in the form DSNAME=name(n) (where n is the generation increment), or by defining each data set as a nonVSAM data set with a generation data set name, as follows:

```
name.GnnnnVnn
```

If the name is found to match the name of a previously defined generation data group, the nonVSAM data set is associated with the generation data group as a generation data set. See "Defining a NonVSAM Data Set" for a description of how to define a nonVSAM entry.

When you create a generation data group it can be cataloged in a VSAM catalog. The model DSCB must exist on the generation data group's catalog volume. To create a generation data group in an OS CVOL, you can use the BLDG function of the OS/VS utility, IEHPROGM. All JCL specifications available to support generation data groups work on VS2 Release 2 and later releases in a way equivalent to OS/MVT, VS2 Release 1, or VS2 Release 1.6, whether the generation data group is cataloged in a VSAM catalog or in an OS CVOL.

Refer to the chapter "Command Format" for a description of the command format for DEFINE GENERATIONDATAGROUP. An example is also included.

DEFINING A PAGE SPACE

The DEFINE command can be used to define an entry for a page space. A page space is an OS/VS2 system data set that contains pages of virtual storage; the pages are stored into and retrieved from the page space by the Auxiliary Storage Manager. A page space is a nonindexed data set (that is, an entry-sequenced cluster) that is entirely preformatted before it is used. In a system with the Mass Storage System, a page space cannot be defined on a mass-storage volume. A page space must reside on a single volume. A page space cannot be opened as a user data set.

A page space has a maximum usable size equal to 65,535 paging slots (records). See the description for space size declarations (CYLINDERS, RECORDS, and TRACKS for DEFINE PAGESPACE) in the "Command Format" chapter for further details.

You can define a page space in a user catalog, then move the catalog to a new system and establish it as the system's master catalog. When you define a page space in a user catalog, you must also code a STEPCAT or JOBCAT statement to identify and allocate the catalog.

When you issue a DEFINE PAGESPACE command, the system creates an entry in the catalog for the page space, and then preformats the page space. If an error occurs during the preformatting process (for example, an I/O error or an allocation error), the page-space entry in the catalog exists even though no space for it exists. You must issue a DELETE command to remove the page-space entry before you redefine the page space.

Each page space is represented by two entries in the catalog: a cluster entry and a data entry (that is, a page space is conceptually an entry-sequenced cluster.) Both of these entries

should be password protected if the page space is to be password protected. (NOTE: The passwords you specify with the DEFINE PAGESPACE command are put in both the page space's cluster entry and its data entry.) When you define page spaces during system generation (SYSGEN), you should use the ALTER command to add passwords to each entry, since passwords cannot be specified during system generation.

In addition, you should ensure that the catalog containing the page-space entry is password protected. Otherwise, a user can list the catalog's contents and find out each entry's passwords.

A page space is made known to the system as a system data set at system-generation time or through members of a partitioned data set: SYS1.PARMLIB. To be used as a page space, a page space must have been defined in a catalog that is currently being used as the master catalog (that is, a page space cannot be used if its entry is in a user catalog).

A page space's passwords and protection attributes are cataloged in both the page-space entry and its data-component entry. You should specify master, control, update, and read passwords to prevent the page space from being accessed if its volume is moved to a VS1 system. The VS2 system automatically prevents a user's program from accessing a page space or its data component.

Refer to the chapter "Command Format" for a description of the command format of DEFINE PAGESPACE. Examples of how to define page spaces are also included.

BUILDING AN ALTERNATE INDEX

The BLDINDEX command is used to build an alternate index. Before you can build an alternate index, you must define and load its base cluster (see the DEFINE CLUSTER command), then define the alternate index and relate it to the base cluster (see the DEFINE ALTERNATEINDEX command). The base cluster and alternate index are described by entries in the same catalog.

The alternate index's volume and the base cluster's volume must be mounted during the build process. Any volumes identified with the WORKFILES parameter must also be mounted.

An alternate index can be built for a key-sequenced cluster or for an entry-sequenced cluster. The base cluster must be nonempty (that is, its high-used RBA value cannot be zero). An alternate index cannot be built to support an empty or reusable cluster, a relative-record cluster, a catalog, another alternate index, or a nonVSAM data set.

Each of the base cluster's data records must be long enough to contain the alternate key. Each record's alternate-key value must be unique, unless the alternate index was defined with the NONUNIQUEKEY attribute.

Refer to the chapter "Command Format" for a description of the command format for BLDINDEX. Examples of the BLDINDEX command are also included.

HOW AN ALTERNATE INDEX IS BUILT

When an alternate index is built by BLDINDEX processing, Access Method Services opens the base cluster to sequentially read the data records, sorts the information obtained from the data records, and builds the alternate-index records:

1. The base cluster is opened for read-only processing. You can prevent other users from updating the base cluster's records during BLDINDEX processing by including the DISP=OLD parameter in the base cluster's DD statement. If INDATASET is specified, Access Method Services will dynamically allocate the base cluster with DISP=OLD.
2. The base cluster's data records are read and certain information is extracted to form the key-pointer pair:
 - When the base cluster is entry-sequenced, the alternate-key value and the data-record's RBA form the key-pointer pair.
 - When the base cluster is key-sequenced, the alternate-key value and the data-record's prime-key value form the key-pointer pair.

If the base cluster's data records can span control intervals, the alternate key must be in the record's first segment (that is, the first control interval).

3. The key-pointer pairs are sorted in ascending alternate-key order. If your program provides enough virtual storage, Access Method Services performs an internal sort (that is, the sorting of key-pointer pairs takes place entirely within main, or virtual, storage).

You can determine the amount of virtual storage required to sort the records internally by using the following process:

- a. Sort record length = alternate key length + (prime key length (for a key-sequenced data set) or 4 (for an entry-sequenced data set)).
- b. Record sort area size = sort record length x number of records in base cluster, rounded up to the next integer multiple of 2048, or 32,768, whichever is greater.

This means the record sort area size will be a minimum of 32,768, or sort record length x number of records rounded up to the next 2K boundary.

- c. Sort table size = (record sort area size ÷ sort record length) x 4.

The sum of b and c above is the required amount of virtual storage for an internal sort. This amount is in addition to the normal storage requirements for processing an Access Method Services command.

If you do not provide enough virtual storage for an internal sort, or when you specify the EXTERNALSORT parameter, Access Method Services defines and uses two sort workfiles and sorts the key-pointer pairs externally. Access Method Services uses the sort workfiles to contain most of the key-pointer pairs while it sorts some of them in virtual storage. This process is called an external sort. An external-sort workfile is a VSAM entry-sequenced cluster, marked reusable, with space suballocated from a VSAM data space. The minimum amount of virtual storage you will need for an external sort is:

$$32,768 + ((32,768 \div \text{sort record length}) \times 4)$$

The amount of space that Access Method Services requests when defining each sort workfile is calculated as follows:

- a. Sort records per block = 2041 sort record length
- b. Primary space allocation in records =
 (number of records in base cluster sort records per block) + 10
- c. Secondary space allocation in records =
 (Primary space allocation x 0.10) + 10

Both primary and secondary space allocation are requested in records with a fixed-length record size of 2041 bytes and a control interval size of 2408 bytes.

4. When the key-pointer pairs have been sorted into ascending alternate-key order, Access Method Services builds an alternate-index record for each key-pointer pair. If the alternate index is defined with the NONUNIQUEKEY attribute and more than one key-pointer pair have the same alternate-key values, the alternate-index record that is built contains the alternate-key value followed by the pointer values in ascending order. If the alternate index is defined with the UNIQUEKEY attribute, each alternate-key value is unique.

When the record is built, it is written into the alternate index as though it is a data record being loaded into a key-sequenced cluster. Attributes and values that apply to the loading of data records can be specified when the alternate index is defined and are: RECORDSIZE, CONTROLINTERVALSIZE, BUFFERSPACE, FREESPACE, WRITECHECK, SPEED, RECOVERY, REPLICATE, and IMBED.

5. When all alternate-index records have been built and loaded into the alternate index, both it and its base cluster are closed. Steps 1 through 4 are repeated for each alternate

index that is specified with the **OUTFILE** or **OUTDATASET** parameter. When all alternate indexes have been built, if external sort workfiles have been defined, they are deleted. Access Method Services finishes processing and issues messages that indicate the results of the processing.

DD STATEMENTS THAT DESCRIBE THE SORT WORKFILES

You can identify VSAM data space that the sort routine can use by specifying two dnames with the **WORKFILES** parameter and supplying two DD statements that describe the workfiles to be defined by Access Method Services. Each workfile DD statement should be coded:

```
//ddname DD   DSNAME=dsname,VOL=SER=volser,  
//           UNIT=devtype,DISP=OLD,AMP='AMORG'
```

ddname

As specified in the **WORKFILES** parameter. If you do not specify the **WORKFILES** parameter and you intend to provide VSAM data space for external-sort workfiles, you identify the workfile DD statements with the names **IDCUT1** and **IDCUT2**.

dsname

A data set name. The OS/VS Scheduler generates a data set name for the workfile if none is provided.

VOL=SER=volser

Required. Identifies the volume owned by the **STPCAT**, **JOBCAT**, or master catalog in which the workfile is to be cataloged. The workfile's space is suballocated from the volume's VSAM data space. You can specify a maximum of five volumes for each workfile. See the section "How an Alternate Index is Built" for a description of how to calculate the amount of space to be suballocated for each sort workfile. If your **BLDINDEX** job requires external-sort workfiles, this space must be available on the volume(s) identified by **volser** or your job will fail.

UNIT=devtype

Type of direct-access device on which the volume is mounted. You can specify a generic device type (that is, 3330, 2314, etc.) or a unit address (that is, 121, 248, etc.). You cannot specify **SYSDA**.

DISP=OLD

Required.

AMP='AMORG'

Required.

If **BLDINDEX** is being used interactively in a TSO environment, these sort workfile DD statements must be contained in a logon procedure.

ORDER OF CATALOG USE: BLDINDEX

The base cluster and alternate-index entries are in the same catalog, because the catalog owns the volume or volumes that contains the cluster and its alternate index.

The order in which catalogs are selected to determine the catalog to contain the external-sort workfile entries (that is, entries that describe each workfile as an entry-sequenced cluster in a suballocated VSAM data space) is:

1. If a catalog is specified with the **CATALOG** parameter, that catalog is selected to contain workfile entries. Otherwise,

2. The user catalog specified in the current job step (STEPCAT) or, if none is specified for the job step, the user catalog specified for the current job (JOB CAT) is selected to contain the workfile entries. If more than one catalog is specified for the job step or job, the first job step or job catalog is selected to contain the workfile entries. Otherwise,
3. If the entries (data set name on the DD statement) are identified with qualified entrynames and the first qualifier is the same as:
 - the name of a user catalog, or
 - the alias of a user catalog,the user catalog so identified is selected to contain the workfile entries. Otherwise,
4. The master catalog is selected to contain the workfile entries.

MODIFYING CATALOG INFORMATION

The ALTER command is used to alter attributes in catalog entries and to rename members of nonVSAM partitioned data sets. To alter an entry, you need to supply its name and the attributes to be altered.

Altering an entry does not normally require that the entry's volume be mounted, because the entry's use of space and the availability of space in the volume's data spaces can be determined by examining the catalog. The entry's volume must be mounted whenever the volume's VTOC must be consulted or modified, such as when a unique component or nonVSAM data set is renamed. The entry's volume must also be mounted when the entry is cataloged in a recoverable catalog. A JCL DD statement can be used to cause the data set or volume to be allocated. In OS/VS2 MVS, a data set or volume can be dynamically allocated by specifying the data-set name or volume serial number if no DD statement is supplied. The volume must be mounted as permanently resident or reserved.

Key-sequenced VSAM data sets consist of three components: an index component, a data component, and the cluster entry itself. Most attributes of the VSAM data set are associated with its data or index components. Certain attributes such as retention period, owner ID, and cluster protection are associated with the cluster entry itself. If you use the cluster name in the ALTER command, only these attributes will be changed. If you use the cluster name in the ALTER command to alter any attribute not directly associated with the cluster entry, Access Method Services will terminate the ALTER request and issue an error message. The converse is also true; if you specify a data or index component name and attempt to alter an attribute directly associated with a cluster entry, Access Method Services will terminate your ALTER command and issue an error message. You must specify the explicit data or index component name in order to alter any of the remaining attributes such as shareoptions, bufferspace, writecheck, the data or index protection attributes, etc. It is recommended, therefore, that you specify data and index component names at DEFINE time in order to facilitate the use of the ALTER command. Otherwise, you will have to use the long system-generated names. Refer to Figure 8 in the chapter "Command Format" to determine which values or attributes you may alter for a particular entry type.

You cannot ALTER any attributes of a user catalog except those directly associated with the cluster name. Data and index component names specified for a user catalog at DEFINE time are meaningless, because the system will always generate standard names. If you specify these system-generated names for the data or index components in the ALTER command, VSAM will return a catalog return code of "8," claiming the entry is not there. Additionally, if you alter the expiration date to an invalid date, such as "99999," Access Method Services will return a condition code of "0" even though the expiration date remains unchanged.

Refer to the chapter "Command Format" for a description of the command format for ALTER. Examples of the ALTER command are also included.

ORDER OF CATALOG USE: ALTER

The order in which catalogs are searched when an entry is to be located to be altered is:

1. If a catalog is specified in the CATALOG parameter, only that catalog is searched. If the entry is not found, a "no entry found" error is returned to the user.
2. Any user catalog(s) specified in the current job step (with a STEPCAT DD statement) is searched. If more than one catalog is specified for the job step, the catalogs are searched in order of concatenation. If the entry is found, no other catalog is searched.

If a STEPCAT catalog is specified and the entry is not found, the JOBCAT catalog is not searched. The catalog search continues with step 3 below.

If no STEPCAT catalog is specified for the job step, and a user catalog is specified for the current job (with a JOBCAT DD statement), the JOBCAT catalog(s) is searched. If more than one catalog is specified for the job, the catalogs are searched in order of concatenation. If the entry is found, no other catalog is searched. Otherwise,

3. If the entry is identified with a qualified entryname and its first qualifier is the same as:
 - the name of a user catalog, or
 - the alias of a user catalog, or
 - the alias of a control volume,the user catalog or control volume so identified is searched. If the entry is found, no other catalog is searched. Otherwise,
4. The master catalog is searched. If the entry is not found, a "no entry found" error is returned to the user.

GENERIC NAMES AND ALTER

To alter entries with qualified names—for example, PAYROLL.74.MAY—you can identify the entry with its full name (that is, all qualifiers) or with all qualifiers but one. The unspecified qualifier is indicated with an asterisk (for example, PAYROLL.*.MAY). This kind of shortened name, PAYROLL.*.MAY, is called a generic name. When you specify the full entryname, only the entry so identified is modified. When you specify all qualifiers of the entryname but one, the entries whose entrynames match the supplied qualifiers are altered, if they contain the type of information specified with the ALTER command's parameters.

For example, when you specify PAYROLL.*, the entries that might be altered are those whose entrynames contain two qualifiers, the first qualifier being PAYROLL. When you specify PAYROLL.*.MAY, the entries that might be altered are those whose entrynames contain three qualifiers, the first being PAYROLL and the third and last being MAY.

When you identify a catalog (that is, when you specify the ALTER command's CATALOG parameter), Access Method Services searches only the specified catalog for generically named entries. If no catalog is specified, the catalog is selected as indicated in "Order of Catalog Use: ALTER" and is searched for generic-named entries.

You must always specify the first qualifier of a qualified entryname. You can substitute an asterisk for no more than one of the other qualifiers.

You must identify all qualifiers in the qualified name. For example, PAYROLL.* cannot be used to identify a qualified name that contains three or more qualifiers, even though its first qualifier is PAYROLL. PAYROLL.74.MAY.* cannot be used to identify a qualified name that has more or fewer than four qualifiers.

Renaming Generically-Named Entries

You can use generic names to rename a group of cataloged objects. To do this, you specify both the entryname and the new name as generic names. For example, if each entryname identified with the generic name A.*.B is to be renamed with the generic name A.*.C, all entrynames that have A as the first qualifier and B as the third and last qualifier are renamed. The new name has A as the first qualifier and C as the third and last qualifier:

Old Name	New Name
A.1.B	A.1.C
A.2.B	A.2.C
A.3.B	A.3.C

If each entryname identified with the generic name A.B.*.D is to be renamed, all entrynames that have A and B as the first and second qualifiers, and D as the fourth and last qualifier, are renamed. If the new generic name is C.*.DATA, the entry names are renamed as follows:

Old Name	New Name
A.B.1.D	C.1.DATA
A.B.2.D	C.2.DATA
A.B.3.D	C.3.DATA

VSAM VOLUME RECOVERY FUNCTION

The VSAM Volume Recovery function removes all VSAM data spaces and resets volume ownership for a volume that cannot be located with its catalog entry (that is, a system failure or I/O error might have damaged the volume entry). The volume's format-4 DSCB (in the VTOC) is reset to remove its ownership from the VSAM catalog. NonVSAM data sets on the volume are not affected. The VSAM catalog that owns the volume is not accessed or modified—the (damaged) volume entry is unchanged.

The VSAM Volume-Recovery function can be used to remove a user catalog from a volume without first deleting each of the catalog's objects (that is, VSAM objects and nonVSAM data sets described with the catalog's entries). The VSAM Volume-Recovery function allows you to remove a VSAM user catalog that has been damaged by a head crash, a standalone DASDI, or a similar accident.

VSAM erases all VSAM data spaces (even when they contain data or a catalog) and rewrites the volume's VTOC to indicate additional free space on the volume. You should use this function only when you cannot access the catalog that owns the volume. You should use the function carefully, in order to prevent unwanted loss of data. When improperly used, this function can contribute to system integrity exposures. See the section "VSAM Volume Cleanup" in the chapter "Data Security and Protection" for further information about the volume-recovery function.

When you use the ALTER command to recover from a damaged volume, you code the command in this format:

```
ALTER    entry[/password]
         FILE(dname)
         REMOVEVOLUMES(volser[b volser...])
```

entryname[/password]
names the master catalog. If the master catalog is password protected, then its master password must also be supplied.

FILE(dname)
specifies the name of a DD statement that describes a volume to be reset. If more than one volume is to be reset, they must be of the same device type. Concatenated DD statements are not allowed. This parameter is required.

REMOVEVOLUMES(volser[b volser...])
identifies the volume(s) on which all VSAM data spaces are to be removed. VSAM ownership of the volume is also relinquished. Volumes owned by the master catalog identified in entryname cannot be specified.

DISPLAYING CATALOG INFORMATION

The LISTCAT command is used to list entries from a catalog. The entries listed can be selected by name or entry type, and the fields to be listed for each entry can additionally be selected. See "Appendix B. Interpreting LISTCAT Output" for an explanation of the output produced as a result of the LISTCAT command.

If entries to be listed are selected by name, the name(s) can be specified in its entirety, can be specified as a generic name, or can be specified with the LEVEL parameter.

Entries are specified by generic name by supplying all but one qualifier of the name. The qualifier omitted is indicated by an asterisk (*). The first qualifier cannot be omitted.

When you specify ENTRIES (A.*), all two-qualifier entrnames that have A for the first qualifier are selected to be listed. When you specify ENTRIES(A.*.B), all three-qualifier entrnames that have A as the first qualifier and B as the last qualifier are selected to be listed. You can further limit the printing by identifying certain entry types to be printed (that is, by specifying NONVSAM, CLUSTER, etc.).

When you specify LEVEL(A.*.B), all entrnames that have A as the first qualifier and B as the third qualifier are selected to be listed—and some of the selected entrnames might have four or more qualifiers (each must have at least three qualifiers and satisfy the A.*.B selection criteria).

When you specify LEVEL(A), all entrnames that have A as the first qualifier (regardless of the number of qualifiers) are selected to be listed.

You cannot specify an "*" as the last qualifier when you use the LEVEL parameter (note that when you specify LEVEL(A), more entries might be listed than when you specify ENTRIES(A.*), even though both ways appear, at first glance, to be identical). If you specify LEVEL(A.*), the LISTCAT operation terminates with an error message.

Restriction for use of REPRO and LISTCAT in same job step:
LISTCAT cannot run in a job step where the catalog is empty when it is opened. To ensure that the LISTCAT correctly reflects the contents of the catalog, it is recommended that LISTCAT be run as a separate job step.

Refer to the chapter "Command Format" for a description of the command format for LISTCAT. Examples are also included.

ORDER OF CATALOG USE: LISTCAT

When the ENTRIES or LEVEL parameter is not specified, or when the command is not executed through TSO, the order in which catalogs are searched when entries are to be listed using the LISTCAT command is:

1. If a catalog is specified in the CATALOG parameter, only that catalog is listed. Otherwise,
2. The first user catalog specified in the current job step (STPCAT) or, if none is specified, the first user catalog specified in the current job (JOB CAT) is listed. Otherwise,
3. If no user catalog is specified in the current job step or job, the master catalog is listed.

When the ENTRIES or LEVEL parameter is specified, or when the command is executed through TSO, the order in which catalogs are searched when entries are to be listed using the LISTCAT command is:

1. If a catalog is specified in the CATALOG parameter, only that catalog is searched. If the entry is not found, a "no entry found" error is returned to the user. Otherwise,
2. Any user catalog(s) specified in the current job step (STEPCAT) or, if none is specified for the job step, any user catalog(s) specified for the current job (JOB CAT). If more than one catalog is specified for the job step or job, the job step or job catalogs are searched in order of concatenation. If the entry is found, no other catalog is searched. Otherwise,
3. If the entry is not found, and the entry's name is a qualified name, and the first qualifier (that is, the first one to eight characters before a period) is the same as:
 - the name of a user catalog, or
 - the alias of a user catalog, or
 - the alias of a control volume,that user catalog or control volume is searched. If the entry is found, no other catalog is searched. Otherwise,
4. The master catalog is searched. If the entry is not found, a "no entry found" error is returned to the user.

DELETING CATALOG ENTRIES

The DELETE command is used to delete entries from a catalog or to delete members of a nonVSAM partitioned data set. When the entry represents an object that contains space in a VSAM data space, the object's space is made available for other VSAM objects.

Deleting an entry does not normally require that the entry's volume be mounted, because the entry's use of space in the volume's data spaces can be determined by examining the catalog. The entry's volume must be mounted whenever the volume's VTOC must be consulted or modified, such as when:

- A data space, a unique component, or a user catalog is to be deleted.
- A cluster or alternate index is to be deleted and erased.
- A nonVSAM data set is to be deleted and scratched.

If the entry is cataloged in a recoverable catalog, its prime catalog recovery area volume must be mounted so that the CRA can be updated to reflect the entry's deletion.

For multivolume VSAM data sets, each volume must be mounted to reflect the entry's deletion from the volume's catalog entry. See "Restoring Catalog Entries After System Failure" for a description of catalog recovery area contents.

JCL can be used to cause a data set or volume to be allocated. A data set or volume can be dynamically allocated by specifying the data set name or volume serial number if no DD statement is supplied. Any volume to be dynamically allocated must be mounted as permanently resident or reserved.

Entries can be deleted from more than one catalog with a single DELETE command. You do this by specifying many entry names and not coding the CATALOG parameter. Use concatenated JOBCAT or STEPCAT DD statements to identify the catalogs that contain the entries.

Note: When two catalogs in your system might contain the same entryname, you should issue separate DELETE commands that reference each catalog.

The kinds of catalog entries you can delete are:

- ALIAS—The alias entry is deleted.

You can delete the alias of a user catalog and redefine it in the same step, but you cannot use it in that step. A TSO session is considered one step so you would have to log off and log on to use an alias of a user catalog that you had deleted and redefined.

Password required: the catalog's update- (or higher-level) password.

- ALTERNATEINDEX—The alternate-index entry, its data component's entry, and its index component's entry are deleted. Space allocated to the alternate index's components is made available for other VSAM objects. Each path entry that relates the alternate index to its base cluster is also deleted.

When ERASE is coded, or when the alternate index has the ERASE attribute, the volume containing the alternate index's data component must be mounted.

The alternate index's components cannot be deleted separately. You must delete the alternate index as a whole. When the alternate index's components reside in unique data spaces, the volumes containing these components must be mounted.

If the alternate index is cataloged in a recoverable catalog, the prime catalog recovery area volume as well as other volumes for the alternate index must be mounted. The prime volume is the first index volume of the base cluster over which the alternate index is defined if the cluster is a KSDS. Otherwise, it is the data volume of the cluster. Its volume serial number is printed by Access Method Services when the alternate index is defined.

Password required: the alternate index's master password or its catalog's master password.

- **CLUSTER**—The cluster entry, its data component's entry, and, if the cluster is key-sequenced, its index component's entry are deleted. Space allocated to the cluster's components is made available to other VSAM components.

Alternate indexes that are associated with the cluster are deleted. The cluster's upgrade-set entry, if it exists, is also deleted. Path entries that relate the cluster to its alternate indexes are also deleted, as well as path entries that have been defined for the cluster alone.

Space allocated to the alternate index's components is made available to other VSAM objects. When the alternate index's components reside in unique data spaces, the volumes containing these components must be mounted.

When ERASE is coded, or when the cluster has the ERASE attribute, the volume containing the cluster's data component must be mounted.

The cluster's components cannot be deleted separately. You must delete the cluster as a whole. When the cluster's components reside in unique data spaces, the volumes containing these components must be mounted; then data spaces are automatically deleted.

When the cluster is cataloged in a recoverable catalog, the prime catalog recovery area of the volume as well as all other volumes of the cluster must be mounted. The prime volume is the first volume containing the index component of the cluster if the cluster is a KSDS. Otherwise, it is the first data volume of the cluster. Its volume serial number is printed by Access Method Services when the cluster is defined.

Password required: the cluster's master password or its catalog's master password.

- **GENERATIONDATAGROUP**—The generation data group (GDG) entry is deleted. The generation data group must be empty (that is, it cannot contain generation data sets) unless you specify FORCE.

Password required: the catalog's update- (or higher-level) password.

Note: On MSS volumes, a GDG data set may not be scratched when it is automatically uncataloged, even if the SCRATCH option has been specified.

- **MASTERCATALOG**—In a VS2 system, the VSAM master catalog is the system's primary catalog. You are not allowed to delete your system's master catalog.

- NONVSAM**—The nonVSAM entry is deleted. When SCRATCH is coded, the data-set's DSCB is removed from its volume's VTOC—the volume containing the nonVSAM data set must be mounted.

All alias entries that point to the nonVSAM entry are also deleted.

When the nonVSAM data set is partitioned, you can delete one of its members by specifying the entryname as

```
pdaname(membername)
```

Password required: the catalog's update- (or higher-level) password.
- PAGESPACE**—The page-space entry and its data-component entry are deleted. Space allocated to the page-space is made available to other VSAM objects.

The page space being deleted must be inactive when it is deleted, and all other page spaces on the same volume must also be inactive.

Password required: The page space's master password or the catalog's master password.
- PATH**—The path entry is deleted. Entries for the path's alternate index and base cluster are undisturbed.

Password required: The path's master password or its catalog's master password.
- SPACE**—All empty data spaces on the volume are deleted. The DSCB entry that describes each deleted VSAM data space in the volume's VTOC is removed. The space is made available to other OS/VS2 system objects.

When all VSAM data spaces on the volume have been deleted, or when you specify FORCE, its volume entry in the catalog is deleted. The volume ownership indicator in the volume's VTOC is reset so that a VSAM catalog no longer owns the volume.

When the last data space on a volume is deleted and the volume is not a candidate volume, the volume's entry is also deleted from the catalog. The effect of deleting the volume entry is that the volume is no longer owned by the catalog, that is, another catalog is free to allocate space on the volume.

The volume must be mounted.

You cannot specify the PURGE parameter when you delete VSAM data spaces.

When you identify data spaces to be deleted, you cannot identify any other type of cataloged object to be deleted with the same command. You can, however, follow the DELETE SPACE command with other Access Method Services commands.

Password required: the catalog's update (or higher-level) password. When you also specify FORCE, you must supply the catalog's master password.
- USERCATALOG**—The catalog's self-describing entries and its user-catalog entry in the master catalog are deleted. The DSCB entry that describes the user-catalog's data space is removed from the volume's VTOC. The volume-ownership indicator in the volume's VTOC is reset so that the volume is not owned by a VSAM catalog.

The catalog must be empty—it cannot contain any catalog entries, unless you specify FORCE.

All alias entries in the master catalog that point to the user-catalog connector entry are also deleted.

You cannot code the ERASE or CATALOG parameters when you delete a user catalog.

The catalog's volume must be mounted. If the volume is unavailable or inaccessible, DELETE cannot be used. You should use EXPORT DISCONNECT instead.

You can delete a user catalog and redefine it in the same step, but you cannot use the redefined catalog in that same step. A TSO session is considered one step so you would have to log off and then log on to use a catalog that you have deleted and redefined.

If there is any possibility of other users accessing a user catalog while it is being deleted, you should serialize the job or job step that deletes the catalog by providing a JOBCAT or STEPCAT DD statement specifying the catalog name as dsname and DISP=OLD.

Password required: the user catalog's master password is specified with its entryname.

You cannot delete a data or index entry separately from its associated cluster, alternate index, or user catalog entry.

If the only object on a VSAM volume is a cluster or alternate index that resides in a unique data space, deleting its entry does not cause VSAM volume ownership to be relinquished. The volume entry remains in the catalog in a candidate status. You must issue a DELETE SPACE command to cause VSAM to release ownership of the volume.

Refer to the chapter "Command Format" for a description of the command format for DELETE. Examples of deleting VSAM and nonVSAM data sets are also included.

ORDER OF CATALOG USE: DELETE

The order in which catalogs are searched when an entry is to be located to be deleted is:

1. If a catalog is specified in the CATALOG parameter, only that catalog is searched. If the entry is not found, a "no entry found" error is returned to the user.
2. Any user catalog(s) specified in the current job step (with a STEPCAT DD statement) is searched. If more than one catalog is specified for the job step, the catalogs are searched in order of concatenation. If the entry is found, no other catalog is searched.

If a STEPCAT catalog is specified and the entry is not found, the JOBCAT catalog is not searched. The catalog search continues with step 3 below.

If no STEPCAT catalog is specified for the job step, and a user catalog is specified for the current job (with a JOBCAT DD statement), the JOBCAT catalog(s) is searched. If more than one catalog is specified for the job, the catalogs are searched in order of concatenation. If the entry is found, no other catalog is searched. Otherwise,

3. If the entry is identified with a qualified entryname and its first qualifier is the same as:
 - the name of a user catalog, or

- the alias of a user catalog, or
 - the alias of a control volume,
- the user catalog or control volume so identified is searched. If the entry is found, no other catalog is searched. Otherwise,
4. If the entry is not found, the master catalog is searched. If the entry is not found in the master catalog, a "no entry found" error is returned to the user.

GENERIC NAMES AND DELETE

To delete entries with qualified names—for example, PAYROLL.74.MAY—you can identify the entry with its full name (that is, all qualifiers) or with all qualifiers except one. The unspecified qualifier is indicated with an asterisk (for example, PAYROLL.*.MAY). This kind of shortened name, PAYROLL.*.MAY, is called a generic name. When you specify the full entryname, only the entry so identified is modified. When you specify all qualifiers of the entryname except one, the entries whose entrynames match the supplied qualifiers are deleted, if they satisfy other parameters specified with the DELETE command's parameters.

For example, when you specify PAYROLL.*, the entries that might be deleted are those whose entrynames contain two qualifiers, the first qualifier being PAYROLL. When you specify PAYROLL.*.MAY, the entries that might be deleted are those whose entrynames contain three qualifiers, the first being PAYROLL and the third and last being MAY.

When you identify a catalog (that is, when you specify the DELETE command's CATALOG parameter), Access Method Services searches only the specified catalog for generically named entries. If no catalog is specified, the catalog is selected as indicated in "Order of Catalog Use: DELETE" and is searched for generic-named entries.

You must always specify the first qualifier of a qualified entryname. You can substitute an asterisk for no more than one of the other qualifiers.

You must identify all qualifiers in the qualified name. For example, PAYROLL.* cannot identify a qualified name that contains three or more qualifiers, even though its first qualifier is PAYROLL. PAYROLL.74.MAY.* cannot identify a qualified name that has more or fewer than four qualifiers.

MOVING ENTRIES

The EXPORT and IMPORT commands allow you to create backup copies of alternate indexes and clusters, to transport user catalogs, alternate indexes and clusters from one system to another, and to prevent (and subsequently, with the IMPORT command, to allow) usage of a user catalog.

When a user catalog is exported, its catalog-connector entry is removed from the master catalog. When the user catalog is subsequently imported to a new system, a catalog-connector entry is created for it in the new system's master catalog and the user catalog is physically transported to the new system. The user-catalog volume itself need not be mounted either to export or import the user catalog. When a user catalog is exported, it is not copied; the user catalog remains on its volume in its original form. The user catalog and its cataloged objects are unavailable to any user until the user catalog is connected (with the IMPORT CONNECT command) to a master catalog.

You do not have to issue the EXPORT command to allow a user catalog to be moved to another system. If the catalog is on a demountable volume, you can physically move the volume to another system. You must use an IMPORT CONNECT command to build a catalog-connector entry for the user catalog in the new system's master catalog. When you issue the EXPORT-DISCONNECT command, VSAM removes the user-catalog's connector entry.

Note: If you issued the EXPORT DISCONNECT command and your user catalog is protected by the RACF Program Product, the RACF profile will not be automatically deleted. It is your responsibility to remove the profile with RACF utility DELDSD. For more information, see the publication OS/VS2 Resource Access Control Facility (RACF) Command Language Reference.

When a cluster or alternate index is exported, its catalog entry is copied to a movable volume along with the object's contents. The entries and components are subsequently copied into a new system.

Exportation of a cluster or alternate index is either permanent or temporary. In permanent exportation, the catalog record is deleted and storage space is freed in the original system. In temporary exportation, both the sending and receiving systems retain a copy of the object, but the copy in the original system is marked to indicate that there is a copy elsewhere. If you want to export permanently a base cluster and its associated alternate indexes, the alternate indexes must be exported first since deletion of a base cluster causes deletion of any alternate indexes defined over it. A base cluster must always be imported prior to importing any of its alternate indexes.

When a cluster or alternate index is exported, its catalog information (that is, the attributes specified with the DEFINE command—except the catalog's name—and modified with subsequent ALTER commands) is moved with the cluster's data records. However, the statistics kept in the catalog entry are lost (that is, the statistics are not available when the entry is subsequently imported). After the import the statistics stored in the catalog entry are the results of the import only. You cannot export an empty cluster or alternate index.

The permanent exportation of a RACF-protected cluster or alternate index causes the deletion of the data-set profiles. A temporary exportation does not. Also, when a RACF-protected data set is exported, the RACF indicators are moved to the portable data set, but the profile is not.

When a previously exported RACF-protected object is imported, you have an option to reuse the old profiles or to establish new ones. The IMPORT command in the chapter "Command Format" will give you more information on RACF options when importing a VSAM object.

Exporting or importing a RACF-protected catalog will cause no change in the status of the RACF indicators in the catalog or the profiles associated with the catalog.

The portable copy of a cluster or alternate index is a variable-blocked, spanned, sequential data set. Each record of a relative-record data set contains an 8-byte header. Each record of all other types of VSAM data sets contains a 4-byte header. The maximum logical record length of the portable data set is based on the maximum record size of the VSAM data set to be exported. Access Method Services uses the value you specified via the maximum subparameter of the RECORDSIZE parameter of the DEFINE CLUSTER or ALTERNATEINDEX command when the data set was defined. The resulting maximum logical record length of the portable data set is as specified below:

- For relative-record data sets, the greater of 276, or VSAM maximum record size + 8
- For all other types of VSAM data sets, the greater of 272, or VSAM maximum record size + 4

The resulting maximum logical record length of the portable data set is limited to 32,760 bytes (Access Method Services does not support the DCB parameter, LRECL=X). Therefore, the largest VSAM maximum record size which can be handled by EXPORT is:

- 32,752 for relative-record data sets
- 32,756 for all other types of VSAM data sets

If the data set to be exported has been defined with a maximum record size greater than the appropriate length shown above, the EXPORT command terminates with an error message.

You may specify, with the DCB parameter of the DD statement for the portable data set, a blocksize other than the default of 2048 bytes.

Refer to the chapter "Command Format" for a description of the command formats for the EXPORT and IMPORT commands. Examples showing how data sets are exported and imported are also included.

EXPORTING AN ENTRY

The EXPORT command is used in conjunction with the IMPORT command to move alternate indexes, clusters, and user catalogs from one system to another, to provide backup copies of clusters and alternate indexes, and to sever the relationship between a user catalog and the master catalog. The EXPORT command cannot be used to move or provide a backup copy of a master catalog, a nonVSAM data set, data space, generation data group, page space, or path. Access Method Services automatically exports along with a cluster or alternate index any paths defined over it. The paths are redefined when the object is subsequently imported.

If a user catalog is to be exported, the catalog is merely disconnected from the master catalog. The user catalog can then be physically (or logically) moved to another system. Even if the catalog is not moved to another system, it is temporarily unavailable to all users, and its aliases are deleted. The person who is responsible for the catalog can issue IMPORT CONNECT to make the catalog available again.

You should take care to list the user-catalog's aliases before you issue the EXPORT command. You can use the LISTCAT command to list the aliases:

```
LISTCAT -  
  CATALOG(mastercatname/password)-  
  ENTRIES(usercatname)-  
  ALL
```

The user-catalog's aliases are listed under the heading "ASSOCIATIONS."

When you use the IMPORT command to make the catalog available to another system, you can issue the DEFINE ALIAS command to reestablish each of the catalog's aliases.

When a VSAM cluster or alternate index that has paths defined over it is exported, the catalog information for those paths is also exported. IMPORT uses this information to reestablish the paths in the target catalog. Therefore, EXPORT can be used to export path entries, but only as parts of a larger structure being exported (that is, a cluster or an alternate index).

IMPORTING AN ENTRY

The IMPORT command is used in conjunction with the EXPORT command to move a cluster, alternate index, or user catalog from one system to another. When a cluster or alternate index is imported, it is automatically reorganized. When you move an object into the new system, you can change some of its attributes (as described with the OBJECTS subparameters).

When a user catalog is moved into a system, you may connect it to the master catalog in the new system, thereby making the user catalog available in the new system. The user catalog might have been disconnected to make it temporarily unavailable to users even though its volume is not moved. The person who is responsible for the catalog can issue IMPORT CONNECT to make it available again.

A cluster, alternate index, or user catalog cannot be moved into a system if its name or the names of any of its components is the same as any name that already exists in the receiving catalog. There are two exceptions to this:

- When IMPORT is used to replace a cluster or alternate index with its backup copy. When the object's backup copy was made, its catalog entry was marked temporary (that is, TEMPORARY was coded when the object was exported). The object's temporary entry is deleted when it is imported; a new entry is built and imported in place of the temporary entry.
- When the cluster or alternate index has been predefined in the receiving catalog. The predefined entry must be for an empty data set. If a duplicate entryname exists for an empty data set, IMPORT assumes that this is a predefined entry whose attributes and space allocation are to be used in place of those in effect when the data set was exported.

If duplicate entrynames are found and the entry in the receiving catalog is not marked as temporary and is not empty, the IMPORT command is terminated.

The first record of the copy to be imported causes an entry to be defined in the receiving system. A check is made for a duplicate name; if a duplicate name exists and is marked temporary, the duplicated name is deleted from the receiving system so that the new entry can be defined. The deletion takes place even if the request for import cannot subsequently be completed.

A check is also made for an empty target data set. If such a data set is present, IMPORT makes no attempt to define a replacement data set. Instead, IMPORT copies the data records from the portable data set into the empty target data set. If you wish to import into an empty target data set, you must specify the INTOEMPTY parameter.

CONVERTING AN OS CATALOG'S ENTRIES TO VSAM CATALOG ENTRIES

When you want to use the entries in an OS catalog with your OS/VS2 system, you can either establish a connector in the master catalog to the OS catalog or you can convert the OS catalog entries to VSAM catalog entries. (Note: The OS catalog is called a control volume or CVOL.) You can convert the OS CVOL entries to VSAM catalog entries using one of two methods:

- Use the DEFINE command to convert some of the entries in an OS CVOL to VSAM catalog entries.
- Use the CNVTCAT command to convert all of the entries in an OS CVOL to VSAM catalog entries in a user or master catalog.

Refer to chapter "Command Format" for a description of the command format for CNVTCAT. Examples are also included.

CONNECTING AN OS CVOL TO THE MASTER CATALOG

When you want to use an OS CVOL in an OS/VS2 system, you must define the CVOL in the system's master catalog. For a description of how this is done, and how OS CVOLs are used in an OS/VS2 system, see OS/VS2 Using OS Catalog Management with the Master Catalog: CVOL Processor.

USING THE DEFINE COMMAND TO CONVERT SOME OF THE CVOL'S ENTRIES

When you issue the DEFINE command to redefine, in a VSAM catalog, objects cataloged in an OS catalog (that is, OS CVOL), you can:

- Selectively catalog entries from the OS catalog
- Convert a large OS catalog to two or more VSAM catalogs, without first splitting the OS catalog

However, this method requires that you issue (that is, keypunch or type in) many DEFINE commands. To lessen the amount of keypunching, you might use the following procedure:

1. Execute the IEHLIST utility program to list the OS catalog. The output of the program is the input to step 2. The IEHLIST program is described in detail in OS/VS2 MVS Utilities.
2. Write and execute:
 - A program to generate DEFINE ALIAS statements for each CVOL pointer entry and alias entry in the OS CVOL listing.
 - A program to generate DEFINE ALIAS statements for each CVOL pointer entry that relates the CVOL pointer names to the appropriate catalog name.
 - A program to generate DEFINE NONVSAM statements for each nonVSAM data set in the OS CVOL listing.
 - A program to generate DEFINE GENERATIONDATAGROUP statements for each generation-index pointer entry in the OS CVOL listing.

The output of step 1 (the IEHLIST listing) is the input to your program. The output of your program is a series of DEFINE ALIAS, DEFINE NONVSAM, and DEFINE GENERATIONDATAGROUP commands.

3. Select the DEFINE statements that represent the entries you want cataloged in one of the VSAM catalogs.
4. Code and execute the appropriate JCL statements to execute the Access Method Services program, IDCAMS, using the selected DEFINE statements from step 3 as SYSIN input statements.
5. Repeat steps 3 and 4 until you have selected and cataloged as many entries as you want in the appropriate VSAM catalogs. CVOL aliases must be put in the VSAM master catalog.

USING THE CNVTCAT COMMAND TO CONVERT ALL OF THE OS CVOL'S ENTRIES TO VSAM CATALOG ENTRIES

The CNVTCAT command is used to convert entries in an OS catalog to entries in an existing master or user catalog. CNVTCAT converts data-set entries, generation data group entries, alias entries, and CVOL-pointer entries to appropriate VSAM catalog entries.

CNVTCAT was designed to be executed only once for an OS catalog. The execution of CNVTCAT is time-consuming for any catalog greater than one hundred entries. Therefore, you should not plan to repeatedly convert the entries of a large OS catalog.

When an OS catalog structure of a system catalog and many OS CVOLs are converted to a VSAM master catalog and many user catalogs, you can use Access Method Services commands to perform the entire operation. You can also merge an OS CVOL's entries into a VSAM catalog, using the CNVTCAT command. However, you cannot easily divide the entries of a large OS catalog among many VSAM user catalogs. When the OS catalog is organized on the basis of qualified data-set names, you can split the OS catalog into two or more OS catalogs using the IEHMOVE function of a selective COPY CATALOG command. You can then convert the entries of the smaller OS catalogs into selected VSAM catalogs. When the OS catalog is not organized by qualified data set name, you can catalog each individual entry using the DEFINE command as illustrated in the section "Using the DEFINE Command to Convert Some of the CVOL's Entries."

OS control volumes (CVOLs) can be interchanged within different operating systems: OS/VS2-1, OS/VS2-2, and OS/VS2-3. When you intend to continue production on another operating system while converting to OS/VS2, you might want to use OS CVOLs because they can be transferred between operating systems.

When you catalog an OS CVOL in the system's master catalog, you can use the CVOL to point to its cataloged entries. You cannot use a STEPCAT or JOBCAT DD statement to identify an OS CVOL, however; an OS CVOL and a VSAM user catalog might reside on the same volume, since the CVOL is considered a nonVSAM data set.

You can issue the DELETE ALIAS command to delete a CVOL's alias in the master catalog. When you issue the DELETE NONVSAM command to uncatalog the OS CVOL, all of its aliases are also deleted.

Figure 4 shows OS catalog-entry types, abbreviations, descriptions, and the VSAM entry types into which OS catalog entries are converted.

Entry Name	Abbreviation	Description	VSAM Entry Type
Alias entry	AE	Contains an alternate name for the high-level qualifier of the data set name	Alias entry.
Control volume pointer entry	CVPE	Connects another control volume (CVOL) to this CVOL	Alias entry for a user catalog in master catalog.
Data-set entry pointer	DSPE	Contains the name and location of a data set	NonVSAM entry.
Generation index pointer entry	GIPE	Points to the lowest index for a generation data group	Generation data group.
Generation data set	G0000V00	Points to a generation data set	NonVSAM entry that is attached to a generation data group.

Figure 4. OS Catalog Entry Types and VSAM Equivalents

When you use the CNVTCAT command, you can cause all the entry types shown in Figure 4 to be converted or all but the control-volume-pointer entries (CVPEs) to be converted. CVPEs are converted to alias entries in the master catalog. If, for example, ABC was the name of a control volume (CVOL) that resided on volume 123456, ABC becomes the alias in the master catalog; you provide the name of the VSAM catalog for which ABC is an alias.

Because CNVTCAT creates aliases for CVPEs, converted OS catalog entries can be found without a STEPCAT or JOBCAT DD statement that identifies the VSAM catalog to be searched. If, for example, ABC.DEF, which was previously contained in an OS catalog, is to be found:

1. Master catalog is searched because no STEPCAT or JOBCAT is specified.
2. ABC is found in the master catalog; it is an alias that is related to a user-catalog entry.
3. User-catalog entry points to the user catalog.
4. ABC.DEF is found as a nonVSAM entry on the user catalog.

Figure 5 shows an OS catalog, SYSCTLG that contains two CVPEs and shows the CVOLs that they point to. Before you can convert the CVPEs, you must be able to provide their volume serial numbers. To learn the volume serial numbers, list the OS catalog to be converted. If SYSCTLG, shown in the top half of Figure 5 is converted, aliases, SYSA and SYSB are created in the master catalog. In a separate conversion, the data sets cataloged in SYSA and SYSB are converted into entries in existing VSAM catalogs. The bottom half of Figure 5 shows the master catalog and two user catalogs after SYSCTLG and SYSA and SYSB have been converted.

The amount of space required in the VSAM catalog depends upon the number and kind of records converted. Each OS catalog entry converts to a nonVSAM entry in a VS catalog. The exception is alias entries. One alias entry in an OS catalog can apply to many data-set names. When the alias is converted, one alias

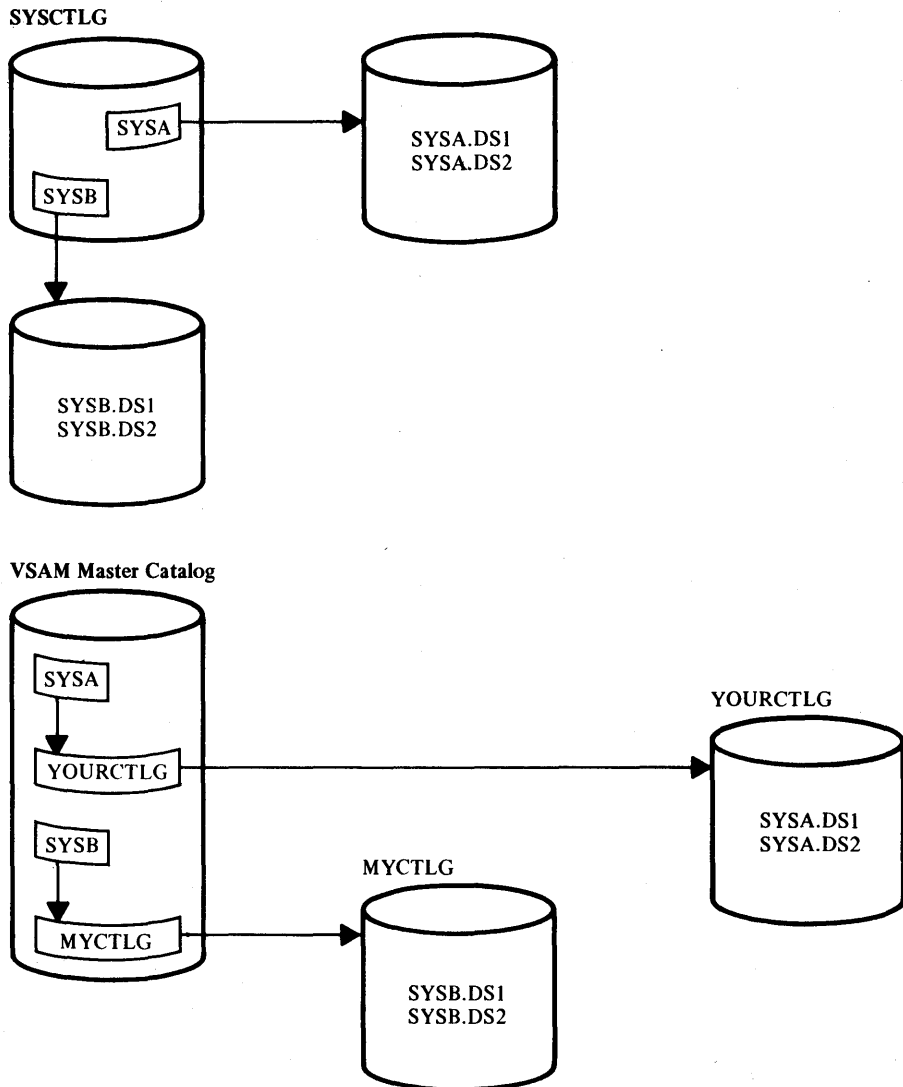


Figure 5. Converting Control Volume Entries

entry is created in the VSAM catalog for each data set that has the high-level qualifier to which the alias applies. For example, an alias, FAKE, exists for REAL. If REAL.A, REAL.B, and REAL.C are converted, a total of six records are required in the VSAM catalog: one record is required for each of the three data sets and one record is required for each of three aliases, one per data set. See "Estimating the Catalog's Space Requirements" to help you estimate the VSAM catalog's size.

Note: Control volumes can be used in a VS2 system without converting entries. See OS/VS2 Using OS Catalog Management with the Master Catalog: CVOL Processor for information about the use of control volumes in a VS2 system.

ORDER OF CATALOG USE: CNVTCAT

The order in which VSAM catalogs are evaluated to select the catalog that receives the converted entries is:

- If a catalog is specified in the CATALOG parameter, that catalog receives the converted entries. Otherwise,
- The user catalog specified for the current job step (STEPCAT) or, if none is specified for the job step, the user catalog specified for the current job (JOB CAT) receives the converted entries. (If more than one catalog is specified for the job step or job, the first catalog in order of concatenation is selected.) Otherwise,
- The master catalog receives the converted entries.

The master catalog always receives alias entries that point to user catalogs.

RESTORING A CLUSTER'S END-OF-FILE VALUES

When a data set is closed, its end-of-file (EOD) and end-of-key-range (EOKR) information is used to upgrade EOD and EOKR information in the data set's cataloged information. If an OS/VS system failure occurs before the data set is closed (that is, before the user's program issues CLOSE or CLOSE(TYPE=T)), the data set's cataloged information is not upgraded. This means that the data set's cataloged information contains possibly obsolete EOD and EOKR information. The data set's real EOD and EOKR indicators are written in the data set, but are not shown in the data set's cataloged information. When the data set is subsequently opened and the user's program attempts to process records EOD or EOKR, a "no record found" error results on a read operation, and a write operation might write records over previously written records.

The VERIFY command is used to compare the end-of-data-set and end-of-key range information as it is stored in a VSAM catalog with the true end-of-file and end-of-key range. If the information in the catalog does not agree with the true end-of-file or end-of-key range, the catalog information is corrected. The VERIFY command can be used following a system failure that caused a component opened for update processing to be improperly closed. Clusters, alternate indexes, and catalogs can be verified. Paths over an alternate index cannot be verified. Paths defined directly over a base cluster can be verified. Although the data and index components of a key-sequenced cluster or alternate index can be verified, the timestamps of the two components will be different following the separate verifies, possibly causing further OPEN errors. Therefore, you should use the cluster or alternate index name as the target of your VERIFY command.

To use the VERIFY command to verify a catalog, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

A user catalog is opened as a data set, when restoring its end-of-file and end-of-key-range information. A JOBCAT or STEPCAT DD statement must therefore be supplied.

Refer to the chapter "Command Format" for a description of the command format for the VERIFY command.

RESTORING CATALOG ENTRIES AFTER SYSTEM FAILURE

When you create a user catalog or a master catalog, you can define it as recoverable. A recoverable catalog has, on each volume it owns, an area that contains a copy of some of the catalog's entries. Each volume's area is called the catalog recovery area or CRA. When a system or device failure occurs that damages the catalog or some of its entries, each volume's catalog recovery area contains the information needed to restore the damaged entries. The contents of a volume's catalog recovery area depend on the types of objects the volume contains. Figure 6 identifies the volume whose catalog recovery area contains a copy of the cataloged object.

Type of Entry	Volume Whose Catalog Recovery Area Contains a Copy of the Catalog Entry
Volume entry	Its own volume
Key-sequenced cluster entry and its data and index entries	The volume that contains the (first part of the) cluster's index component
Alternate-index entry and its data and index entries, when the alternate index is for a key-sequenced cluster	The volume that contains the (first part of the) alternate index's base cluster's index component
Path entry, when the path is related to a key-sequenced cluster	The volume that contains the (first part of the) path's base cluster's index component
Entry-sequenced cluster's entry and its data entry	The volume that contains the (first part of the) cluster's data component
Alternate-index entry and its data and index entries, when the alternate index is for an entry-sequenced cluster	The volume that contains the (first part of the) alternate index's base cluster's data component
Path entry, when the path is related to an entry-sequenced cluster	The volume that contains the (first part of the) path's base cluster's data component
Relative-record cluster entry and its data entry	The volume that contains the (first part of the) cluster's data component
NonVSAM data set	The volume that contains the nonVSAM entry's catalog
Generation Data Group	The volume that contains the entry's catalog
Alias Entry	The volume that contains the catalog in which the nonVSAM association for the alias is defined
User-catalog connector entry in the master catalog	The volume that contains the master catalog
Catalog's self-describing entries	These entries are not duplicated in any catalog recovery area

Figure 6. Catalog Recovery Area Contents

If you should discover that your recoverable catalog is damaged such that its entries are inaccessible, are downlevel, or contain erroneous information, you have the option of choosing one of two different methods to restore your catalog to a usable condition.

- The EXPORTRA/IMPORTRA Method: You would most likely use this method to selectively repair specific catalog entries. Also, remember that reorganization of your catalog and your data is a byproduct of this approach since it involves the movement of data. You may use the following procedure to restore the usability of your catalog.
 1. Issue the LISTCRA command to list and, optionally, compare the catalog recovery area entries (that is, the catalog recovery area that contains copies of the damaged catalog entries).
 2. Issue the EXPORTRA command to obtain a copy of the damaged entries from the catalog recovery area and, when a VSAM data set entry is moved, obtain a copy of the data set's contents.
 3. Clear the damaged volume or reset it so that it is usable by doing one or more of the following:
 - Issue the DELETE SPACE command with the FORCE parameter to remove VSAM data spaces from the volume.
 - Execute the IEHDASDR program with the DUMP and RESTORE statements to restore nonVSAM data sets on the volume. (See OS/VS2 MVS Utilities for details.)
 - Execute the IEHDASDR program with the ANALYZE statement to initialize the volume. (See OS/VS2 MVS Utilities for details.)
 - Issue the ALTER REMOVEVOLUMES command to remove the VSAM catalog's ownership of the volume.
 4. Issue the IMPORTRA command to reload the data copied and moved during step 2 above.
- The RESETCAT Method: If you do not want your data to be moved and if you wish to confine all updating to the catalog (and CRAs), you should consider this approach. RESETCAT does not permit selective reset of specific catalog entries. An entire volume's worth of catalog entries are reset. You would use RESETCAT if a catalog or one or more of its owned volumes become inaccessible. Restore the volume(s) from a backup copy and issue RESETCAT to provide the necessary consistency between the catalog and the restored volume(s).

The LISTCRA, EXPORTRA, IMPORTRA, and RESETCAT commands are described in detail in the sections that follow.

LISTING THE CATALOG RECOVERY AREA'S CONTENTS

If your catalog was created with the RECOVERABLE option, you can determine the damage (if any occurred) that was done to your catalog when a system or hardware failure occurred. You can also determine if a volume or volumes owned by your catalog is out-of-synchronization with the catalog itself (caused, for example, by restoring a volume from a downlevel backup). This is because a recoverable catalog maintains a copy of each catalog entry in a separate part of the volume, called the catalog recovery area (CRA).

When you issue the LISTCRA command with the COMPARE parameter, Access Method Services compares each entry in the CRA to its corresponding entry in the catalog. The comparison is made on a record-by-record, byte-by-byte basis. When a mismatch is encountered within a record, Access Method Services determines the field within the record in which the mismatch exists and prints a message. For certain fields, the message identifies the field(s) which mismatched. (The section "Regaining Access to Data" lists the mismatches detected by LISTCRA.)

When a mismatch is detected, Access Method Services prints both the catalog record and the corresponding CRA record. Asterisks are printed below the specific area which showed a mismatch. All records associated with the mismatched record are also printed (for example, a volume record and its extensions). You can use the LISTCRA output listing to determine which catalog entries are no longer accurate, and to help you code the ENTRIES subparameter of the EXPORTRA command.

To use the LISTCRA command, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

The types of output listing that the LISTCRA command can produce, depending on the optional parameters, are:

- A list of the name and volser of each entry, and the name and volser of each related entry, in the catalog recovery area. The entries are listed in alphameric order by group type (NOCOMPARE and NAME).
- A full dump (that is, hexadecimal and character listing) of each entry and its related entries in the catalog recovery area. The entries are listed in alphameric order by group type (NOCOMPARE and DUMP).
- A list of the name of each catalog entry and its volser whose data does not compare equally with the entry's copy in the catalog recovery area, and an indication of the type of information that mismatches. The mismatched entries are listed in alphameric order by group type. The entries that compare equally are not listed (COMPARE and NAME).
- A full dump (that is, hexadecimal and character listing) of each catalog record whose contents does not compare equally with the record's copy in the catalog recovery area, and asterisks below each byte that mismatches. The mismatched entries are listed in alphameric order by group type. Refer to the section "Regaining Access to Data" in the chapter "Data Security and Protection" for a description of the types of mismatches detected by LISTCRA. The entries that compare equally are not listed (COMPARE and DUMP).
- A full dump (that is, hexadecimal and character listing) of all entries in the catalog recovery area in sequential order as they occur in the CRA (SEQUENTIALDUMP). "Appendix D. Interpreting LISTCRA Output Listings" provides a description of LISTCRA outputs.

Refer to the chapter "Command Format" for a description of the command format of LISTCRA. An example is also included. "Appendix D. Interpreting LISTCRA Output Listings" provides a description of LISTCRA output.

COPYING A CATALOG ENTRY FROM THE CATALOG RECOVERY AREA

When you discover that your catalog is partially or completely damaged due to a system failure or hardware problem, you want to reconstruct the damaged catalog entries so that you can access the cataloged object's data. When you discover that some of the entries in your catalog are not in synchronization with volumes the catalog owns, you want to resynchronize the incorrect entries so they properly reflect the actual status of the volumes. You can rebuild a catalog entry by issuing the IMPORT command when you have recently made a copy of the object with the EXPORT command. The IMPORT command replaces the damaged catalog entry with its copy in the exported file. (See "Moving Entries" for details about the EXPORT and IMPORT commands.)

An exported copy of your cluster might not exist, or the one you have might not be current. The volume that contains your data might not be damaged but, because its catalog entry is damaged or out-of-synchronization, you cannot use the catalog entry to locate and access your data. What you want to do is replace the damaged catalog entry with its undamaged copy in the catalog recovery area and, when the entry is for a VSAM data set, to move the contents of the data set to a volume owned by an undamaged catalog. For a recoverable catalog, you can use the EXPORTRA command to obtain the catalog-entry's copy from the catalog recovery area, then use the IMPORTRA command to replace the damaged entry with its copy.

If an entire VSAM volume becomes unusable, and a backup copy of the volume exists, you may want to consider using RESETCAT rather than EXPORTRA to reset your catalog so that it will correctly access the VSAM data sets on the restored volume. See "Resetting Catalog Entries."

To use the EXPORTRA command, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

When a VSAM data set is recovered using the EXPORTRA and IMPORTRA commands, the result is essentially the same as if the data set were backed up using the EXPORT and IMPORT commands. The differences in the process used to achieve this result are:

- The data set's catalog entry is copied from a catalog recovery area instead of the VSAM catalog.
- Many data sets can be recovered with one issuance of the EXPORTRA and IMPORTRA commands.

The EXPORTRA command copies VSAM catalog entries and the contents of nonempty VSAM data sets to a movable storage device (that is, a magnetic tape or demountable disk pack). Other capabilities of the EXPORTRA command that are not available as functions of the EXPORT command are:

- All entries in a catalog recovery area can be copied, including VSAM clusters and alternate index entries, VSAM user catalog connectors and their aliases (from a master catalog CRA only), GDG base entries, and nonVSAM entries and their aliases. The connection between a nonVSAM entry and its GDG base entry is mentioned in the copy. Page spaces are not copied.
- You can copy the entries for an empty cluster or alternate index.
- You can obtain a copy of each entry in the catalog recovery area by issuing the EXPORTRA command once. Subsequently, you can replace all entries exported with the EXPORTRA command by issuing the IMPORTRA command once.

- You can obtain the copy of one or more entries, as you specify, without obtaining the rest of the catalog recovery area.
- When the catalog owns more than one volume, you can obtain the copy of each entry from one volume's catalog recovery area without obtaining entries from the catalog recovery areas of other volumes.

Use of the EXPORTRA command is limited to a single execution per system at one time only. That is, multiple executions cannot be processed simultaneously.

The EXPORTRA command uses the variable-blocked spanned sequential format (SAM RECFM = VBS) for its output data set. Each record contains an 8-byte header. Each data record of a VSAM relative-record data set contains an additional 4-byte header. EXPORTRA sets the maximum logical-record length of the output data set based on the largest maximum recordsize of all the VSAM data sets being copied. To determine the largest maximum record size, EXPORTRA uses the value specified via the maximum subparameter of the RECORDSIZE parameter of the DEFINE CLUSTER or ALTERNATEINDEX command when each of the data sets to be exported was defined. The resulting maximum logical record length of the output data set is as specified below:

- If a VSAM relative record data set has the largest maximum record size, the greater of 280, or the largest VSAM record size + 12
- If a VSAM key-sequenced or entry-sequenced data set has the largest maximum record size, the greater of 280, or the largest VSAM maximum record size + 8

The resulting maximum logical-record length of the output data set is limited to 32,760 bytes (Access Method Services does not support the DCB parameter LRECL=X). Therefore, the largest VSAM record which can be handled by EXPORTRA is:

- 32,748 for relative record data sets
- 32,752 for all other types of VSAM data sets

If any of the VSAM data sets to be exported has been defined with a maximum record size greater than that shown above, the EXPORTRA command terminates with an error message.

You may specify, with the DCB parameter of the DD statement for the portable data sets, a block size other than the default of 2048 bytes.

If the EXPORTRA command is executed for a RACF-protected VSAM entity, the RACF indicator will be exported on to the portable data set. However, profiles will not be deleted, nor will the profiles be moved to the portable data set. When RACF-protected VSAM entities are imported using IMPORTRA, you have an option to reuse the old profiles or to establish new ones.

The EXPORTRA command requires a DD statement for each volume to be accessed; all volumes must be mounted. EXPORTRA does not dynamically allocate the required volumes.

Refer to the chapter "Command Format" for a description of the command format for EXPORTRA. Examples of EXPORTRA are included with the EXPORTRA command format in the same chapter.

RESTORING THE CATALOG ENTRY THAT WAS OBTAINED USING THE EXPORTRA COMMAND

The IMPORTRA command is used to reestablish in a catalog all those objects that reside in the portable data set created by a previously issued EXPORTRA command. When an existing catalog entry is found with the same entryname as the object in the portable data set, the existing entry is first deleted. The object is redefined in the catalog, using information from the portable data set.

VSAM clusters, their associated data and index components, and any paths over them, alternate indexes, their associated data and index components, and any paths over them, nonVSAM data sets, and generation data group entries are automatically defined in the catalog selected by the user. If a VSAM cluster or alternate index was not empty at the time the EXPORTRA command was issued, its data records were copied to the portable data set together with its VSAM catalog entries. These data records are reloaded into the space occupied by the redefined object. User catalog connector entries (which can be exported only from the system's master catalog) are connected to the master catalog. (Existing user catalog-connector entries with the same entryname as the imported entry are disconnected rather than deleted; the imported user catalog-connector entry is then reconnected.) The aliases of user catalog and nonVSAM entries are also redefined by IMPORTRA.

IMPORTRA requires you to provide a DD statement (specified through the OUTFILE parameter) which identifies a data-set name and the volume serial number of each volume that is to contain the imported VSAM clusters or alternate indexes. You must use concatenated DD statements if the data sets are on different device types. The data-set name you specify is used by Access Method Services for its internal processing during the execution of the IMPORTRA command. The sequence of processing steps is shown below:

1. The cluster or alternate-index catalog information is obtained from the portable data set (created by EXPORTRA).
2. The cluster or alternate index is defined in the catalog.
3. If an existing catalog entry with the same entryname is found, the existing entry is deleted. The cluster or alternate index is then redefined in the catalog.
4. The cluster or alternate index entry in the catalog is renamed using the data-set name you provided in the DD statement identified in your OUTFILE parameter.
5. The cluster or alternate index is opened, loaded with its data records, and closed.
6. The cluster or alternate-index entry is renamed to the original name contained in the portable data set.

If a system failure should occur after step 4 and prior to the successful completion of step 6, the data set may exist in the VSAM catalog under the name provided on the DD statement. To correct the situation, you can delete the cluster or alternate index using the data-set name you provided in the DD statement prior to rerunning the IMPORTRA command.

Refer to the chapter "Command Format" for a description of the command format for IMPORTRA. Examples of EXPORTRA and IMPORTRA are included.

RESETTING CATALOG ENTRIES

When you define a catalog as recoverable, each volume owned by the catalog contains a catalog recovery area (CRA). The CRA contains duplicate information for catalog entries associated with that volume. You can use the RESETCAT command when a recoverable catalog or one or more of its owned volumes becomes inaccessible. You can restore the inaccessible volume(s) from a backup copy and execute the RESETCAT command. The CRAs contain enough information to reset the catalog entries, and VSAM data sets owned by that catalog can again be accessed correctly.

Unlike the EXPORTRA/IMPORTRA command, the RESETCAT command is a one-step operation that enables you to recover your catalog without movement of data. The RESETCAT command does not check or process the data itself, but compares catalog entries with CRA entries and resets as necessary in order to enable you to regain access to the data. You are responsible for ensuring that the data is at the correct level for your use.

If a VSAM volume becomes inaccessible, and a backup copy of the volume is used to restore the volume to a previous level, the volume and the catalog may no longer be synchronized. A list created by the LISTCRA command (with COMPARE option) can indicate mismatches that require RESETCAT command to be run. The RESETCAT command can synchronize the catalog with the volume. After access to the data has been regained, the data sets on that volume can be brought up to the current level by rerunning the jobs that were run after the backup was taken.

If a recoverable catalog becomes unusable, use the LISTCRA command to help analyze the problem (also see "Catalog Recovery"). If you are unable to access your data, restore the catalog volume. Then run the RESETCAT command to synchronize the catalog with its owned volumes. If volumes have been added since the catalog backup was made, RESETCAT can build these entries in the catalog from the volume's CRA. If volumes have been deleted since the last backup, use the DELETE SPACE (FORCE) command to delete the volume's space entries in the catalog and delete the data sets that resided on those volumes that are now marked unusable in the catalog.

RESETCAT uses the entries in the catalog recovery area for each volume you identify to synchronize the catalog with these volumes. The actions taken by RESETCAT to accomplish this synchronization are summarized below:

- If an entry exists in the catalog but not in the CRA, the entry is deleted from the catalog.
- If an entry exists in the CRA but not in the catalog, the entry is inserted into the catalog.
- If, in adding an entry to the catalog a duplicate name is encountered, the entry to be added is renamed. If the data or index component of a unique cluster or alternate index is renamed, the corresponding format-1 DSCB is also renamed. If a nonVSAM entry is renamed, the corresponding format-1 DSCB is not renamed.
- The data space accounting in the volume entry is checked against the Volume Table of Contents (VTOC).
 - If the VTOC contains a format-1 DSCB for VSAM space but the CRA volume entry does not reflect this space, the format-1 DSCB is scratched.
 - If the CRA shows space which is not reflected in a format-1 DSCB, the space is deleted from the CRA volume entry and the VSAM data sets which were shown as being contained within the space are marked unusable.

If the extents in the CRA volume entry do not match those in the associated format-1 DSCB, the CRA is adjusted.

- A space consistency check is performed to ensure that the space claimed by a VSAM data set has the correct extents. If a data or index component claims space that is not allocated to VSAM, the component is marked unusable. If a data or index component claims space which is in conflict with another claim, the component in error is marked unusable and all space not shown as claimed will be returned for suballocation.
- An association check is performed to ensure completeness of the structures; for example, a cluster entry must be correctly associated with its data and index components. When a structure is found to be incorrect, it is deleted and its space returned for suballocation. If a base cluster association is incorrect, the entire structure (base cluster, all alternate indexes, and paths) is deleted. If an index structure is incorrect, only the alternate alternate-index structure and its paths are deleted.

Whenever RESETCAT takes action as indicated above, you are informed with appropriate messages.

If your catalog becomes unusable and no backup copy is available, you can use RESETCAT to recover all of your catalog entries:

If the catalog is a user catalog, remove the catalog connector entry of your catalog from the master catalog via EXPORT DISCONNECT, and define a catalog with the same name on a different volume with the RECOVERABLE attribute. The new catalog can be on a different device type. You may specify a different allocation from that of the original catalog. Volumes owned by the unusable catalog should not be included as owned by the new catalog. The DEFINE operation would flag this as an error condition, because volumes are already owned.

With this new catalog, issue RESETCAT, specifying all volumes owned by the previous catalog (including the unusable catalog's resident volume) for reset. Because the new catalog name is the same as the old catalog name, all entries in the specified CRAs will be added to the new catalog (including volume entries). At the conclusion of RESETCAT processing, the old catalog will have been deleted and the space freed for suballocation.

While the catalog is always updated during RESETCAT processing, the CRA can also be updated under certain circumstances. If some external event such as a power failure were to cause RESETCAT to fail, partial updates to the catalog and CRA(s) may have taken place. Therefore, the catalog and any CRA volumes being reset should be restored before RESETCAT is rerun. It is advisable, therefore, for you to have backup volumes of your catalog and CRA(s) before you use RESETCAT.

Prior to running RESETCAT, you may also wish to execute a LISTVTOC for the volumes to be used in the reset operation (because the VTOC may be changed by RESETCAT). Also, you may wish to execute a LISTCAT command for the catalog to be reset.

After RESETCAT processing has completed, you should examine the messages it has issued. These messages describe specific actions taken by RESETCAT, such as marking a data set unusable or deleting an incorrect VSAM structure. You may also find it advisable to execute a LISTCRA command with the COMPARE option to verify that no further mismatches exist.

Any further action you should take depends on each message issued by RESETCAT. If a data set has been marked unusable, the data that is accessible can be copied using the REPRO command.

Data sets marked unusable can be deleted using the DELETE command. If renaming has taken place, you may wish to alter the name selected by RESETCAT to a more meaningful name using the ALTER command with the NEWNAME option.

Refer to the chapter "Command Format" for a description of the command format of RESETCAT. Examples are also included.

RESETCAT REQUIREMENTS

In planning to use RESETCAT, you should be aware of the following requirements:

- Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.
- The catalog being reset must be capable of being opened (that is, any errors terminating OPEN must be resolved before using RESETCAT), and it must have the RECOVERABLE attribute. It may or may not have valid entries.
- The CRAs must be capable of being opened (that is, any errors terminating OPEN must be resolved before using RESETCAT). Entries not related to the CRA itself may be inaccessible.
- The CRAs must have been created by a recoverable catalog with the same name as the catalog being reset.
- The catalog must be extendable in the event that it becomes enlarged as a result of the reset operation.
- If the master catalog is password protected, the master password of that catalog is required.
- The master catalog may not be reset while it is in use as a master catalog.
- No VSAM data sets cataloged in the catalog being reset can be open.
- You need to use caution when using RESETCAT to recover accessibility of a volume which contains a portion of a multivolume file. Prior to issuing RESETCAT, compatible levels of volumes containing multivolume files should be restored. See "Considerations for Multivolume Data Sets" in this chapter for more details.

WORKFILE SPACE REQUIREMENTS

The RESETCAT command requires a temporary work file for use as temporary storage while processing the command. It is defined by the RESETCAT command in a catalog other than the one being reset and deleted at the end of command processing. The space required is suballocated from VSAM data spaces on the volumes assigned on the DD statement for the WORKFILE parameter. Under normal conditions (no extensions), the amount of space required will be no larger than the resultant catalog. You can determine this by a LISTCAT listing of that catalog.

If the catalog must be extended as a result of RESETCAT processing (this may occur when the catalog is restored at a lower level than its owned volumes), enough data space must be provided to allow for this extension. The space required for each extension is 6603 records, where the record size is 505 bytes. An additional 7 bytes per record will be required; the CI size will be set at 512 bytes.

CONSIDERATIONS FOR MULTIVOLUME DATA SETS

The contents of a volume's CRA depend on the types of objects the volume contains. It is important for you to understand on which CRA the catalog information resides for a particular object. Figure 6 in this chapter identifies by object type the location of the CRA.

The primary CRA contains all of the catalog records necessary to describe the object. Hence, for an entry sequenced data set on two volumes, the volume that contains the first part of the entry-sequenced data set contains all the records that describe the entry entry-sequenced data set (including its allocation on the second volume). The second volume, a secondary CRA for this object, contains information that shows that the entry-sequenced data set is allocated on the second volume. If the second volume has an I/O error that renders it useless and a previous version of that volume is restored, the present catalog information may be erroneous; that is, the catalog may reflect the data set's extent on the second volume, whereas there is no longer an extent on the second volume. Prior to issuing a RESETCAT command, you should restore compatible levels of volumes containing multivolume data sets. RESETCAT would then be issued to reset the catalog to reflect the restored level of all data sets on all reset volumes.

If, in the above example, the second volume was restored and RESETCAT specified only that volume as a reset volume, the entry-sequenced data set may be marked unusable for that volume and the space allocated to it would be either scratched or returned to the catalog for suballocation. The primary description is on the first volume, which was not indicated as one to reset. The description of the data set used would be the description that currently resides in the catalog. If the data set is defined differently on the second volume (for example, extents do not match), the data set is marked unusable for that volume and the allocated space marked free.

For a multivolume entry sequenced data set, a multivolume key-sequenced data set, or an alternate index defined on a volume different from the data set it is based on, minimizing the intersection of different multivolume data sets on a common volume will permit better use of RESETCAT.

When all volumes of a multivolume VSAM data set or structure are not specified in the RESETCAT operation, the extent of checking depends on whether the primary CRA volume is specified for reset. If the primary CRA volume is specified for reset, all information in the catalog is replaced for the data set concerned. For all volumes of the multivolume data set, whether specified or not, the following consistency checks are made by RESETCAT:

- Check the current catalog (if the volume is not specified) or the CRA (if the volume is specified) to ensure that the data set is defined on the volume.
- Check the data set specified on each volume. Was it defined at the same time as the one specified in the primary CRA?
- Check the extents described on the volumes. Are they still allocated to the multivolume VSAM data set?

Although the above checks guarantee that the catalog physically describes a data set correctly, these checks cannot guarantee that the level of data in the data set is at a consistent level. For instance, if a multivolume keyed-sequential data set was defined with the data on one volume and the index on another, the same Define-time would be associated with both. If, over some time, several additions, deletions, and updates were made without causing an extension of the data set, RESETCAT would be

unable to distinguish among different combination of volumes taken from this time period. Since the index contains direct VSAM pointers to the data, an inconsistent combination may cause errors.

If the primary CRA volume is not specified for reset, the scope of checking is limited to volumes specified in the reset. The current catalog is checked to ensure that the current catalog entry describes the part of the data set on the reset volume. Hence, only verification (no reset) occurs for these partial entries. The check ensures that the part of the data set on the reset volume resides on the same physical place as described in the current catalog and is part of the same definition as the data set described in the current catalog. RESETCAT cannot guarantee that the level of data in the data set is at a consistent level among different volumes.

JCL REQUIREMENTS

For the catalogs required during RESETCAT processing, the catalog being reset (indicated by the CATALOG parameter of the RESETCAT command) must appear in a STEPCAT or JOBCAT DD statement. The catalog in which the work file is defined (optionally indicated by the WORKCAT parameter) must appear in a STEPCAT or JOBCAT DD statement. If the WORKCAT parameter is not specified, the work file will be defined in the catalog specified in the first concatenation of the STEPCAT or JOBCAT DD statement.

The relationship of the STEPCAT or JOBCAT DD statement is summarized below:

1. If you want the work file defined in a catalog other than the master catalog, then that catalog must be specified first in the JOBCAT or STEPCAT DD statement concatenation.

```
//STEPCAT DD DSN=workcat,DISP=SHR
//          DD DSN=resetcat,DISP=SHR
```

The above example specifies the work file catalog and reset catalog, respectively, in the JOBCAT or STEPCAT DD statement for case (1).

2. If you want the work file defined in the master catalog, then the master catalog must be specified by name via the WORKCAT parameter and, if specified in the STEPCAT or JOBCAT DD statement, must appear last in the concatenation sequence.

```
//STEPCAT DD DSN=resetcat,DISP=SHR
//          DD DSN=mastcat,DISP=SHR
```

The above example specifies the reset catalog and master catalog, respectively, in the JOBCAT or STEPCAT DD statement for case (2).

Further, the catalog being reset must be used by the RESETCAT command as a data set. A separate DD statement may be used for this catalog. It should specify only the catalog being reset and should not be concatenated to another catalog. For example:

```
//DDCAT DD DSN=catname,DISP=OLD
```

DISP=OLD should be specified to ensure exclusive use of the catalog. If no DD statement is supplied, it will be dynamically allocated.

For CRAs, a single DD statement is required for each volume containing a CRA if CRAFILES parameter of the RESETCAT command is specified. For example:

```
//DDCRA1 DD UNIT=3330,VOL=SER=XYZ,DISP=SHR
//DDCRA2 DD UNIT=2314,VOL=SER=ABC,DISP=SHR
```

Unit affinity may be specified to reduce unit requirements. No two CRAs will be demanded concurrently by RESETCAT.

For the work file, RESETCAT command requires a list of one or more volumes to define a temporary VSAM data set; no more than five volumes may be specified. If no data-set name is provided, a system-generated data set name will be used. The following is an example of a work file DD statement:

```
//FILEW DD DSN=A.WFILE,UNIT=(3330,2),VOL=SER=(X,Y),  
// DISP=OLD,AMP='AMORG'
```


COPYING AND PRINTING

The REPRO and PRINT commands are used to copy and print VSAM and nonVSAM data sets, and to copy VSAM catalogs.

COPYING DATA SETS

You use the REPRO command to do any of the following:

- Copy a VSAM data set into another VSAM data set.
- Copy a sequential data set into another sequential data set.
- Copy an alternate index as though it were a key-sequenced VSAM data set.
- Copy a SAM, ISAM, or VSAM data set (whose records are fixed-length) into an empty VSAM relative-record data set.
- Convert a sequential or indexed-sequential data set into a VSAM data set.
- Convert a VSAM or indexed-sequential data set into a sequential data set.
- Copy a data set—other than a catalog—to reorganize it. Data set reorganization is an automatic feature.
- Copy a nonrecoverable catalog from one volume to another, for example, from a 2314 volume to a 3330 volume.
- Make a backup copy of a catalog.
- Reload a backup copy of a catalog.
- Merge two VSAM data sets.

Throughout the remainder of the REPRO discussion, all of these functions will be referred to as copying.

Refer to the chapter "Command Format" for a description of the command format for REPRO. Several examples are included.

VSAM data sets used as either input or output must be cataloged. Sequential and indexed-sequential data sets need not be cataloged.

If a sequential or indexed-sequential data set is not cataloged, you must include the appropriate volume and unit parameters on your DD statement. You must also supply a minimum set of DCB parameters when the input data set is sequential or indexed-sequential, and/or output data set is sequential. The following table shows the four key parameters.

Parameters	Must Be Supplied by User	Default If Not Supplied
DSORG	IS	PS
RECFM	F, FB, V, VB	U
BLKSIZE	blocksize	None
LRECL	lrecl	BLKSIZE for F or FB BLKSIZE-4 for V or VB

Note that the one parameter that is not supplied by default is BLKSIZE; you must supply this value. The DCB parameter DSORG must be supplied via the DD statement. The DCB parameters RECFM, BLKSIZE, and LRECL can be supplied via the DSCB or header label of a standard labeled tape; otherwise, they must be supplied via the DD statement.

For a variable record length VSAM data set, the VSAM RDF—a 4-byte field defining the record length—is not included in the VSAM record length. Thus, when REPRO attempts to copy a VSAM record whose length is within 4 bytes of MACRECL, a recoverable error occurs and the record is not copied.

Access Method Services does not support records greater than 32,760 bytes for nonVSAM data sets (that is, LRECL=X is not supported). If the logical-record length of a nonVSAM input data set is greater than 32,760 bytes, or, if a VSAM data set defined with a record length greater than 32,760 is to be copied to a sequential data set, your REPRO command will terminate with an error message.

Records in an indexed-sequential data set that have a fixed-length, unblocked format with a relative-key-position (RKP) of zero are preceded by the key string when used as input. Therefore, the records in the output data set must have a record length defined that includes the extended length caused by the key string. Also, to copy "dummy" indexed-sequential records (records with hexadecimal 'FF' in the first byte) you must specify the DUMMY option in the ENVIRONMENT parameter.

Because data is copied as single logical records in either key order or physical order, automatic reorganization takes place. The reorganization can cause any of the following:

- Physical relocation of logical records
- Alteration of a record's physical position within the data set
- Redistribution of free space throughout the data set
- Reconstruction of the VSAM indexes

Figure 7 describes how the records from the input data set are added to the output data set when the output data set is an empty or nonempty entry-sequenced, key-sequenced, sequential data set, or relative-record data set.

	Empty	Non-Empty
Entry-Sequenced/ Sequential	Creates new data set in sequential order.	Adds records in sequential order to the end of the data set. (DISP=MOD must be specified for SAM sequential.)
Key-Sequenced	Creates new data set in key sequence and builds an index.	Merges records by key and updates the index. Records whose key duplicates a key in the output data set are lost, unless you specify the REPLACE option.
Relative-Record	Creates a new data set in relative-record sequence, beginning with 1.	Records from another relative-record data set are merged, keeping their old record numbers. A new record whose number duplicates an existing record number is lost, unless you specify the REPLACE option. Records from any other type of organization cannot be copied into a nonempty relative-record data set.

Figure 7. Adding Records To Various Types of Output Data Sets

Note: The copy operation is terminated if:

- One physical I/O error is encountered while writing to the output data set, or
- A total of four errors are encountered in any combination of the following:
 - Logical error while writing to the output data set
 - Logical error while reading the input data set
 - Physical error while reading the input data set

When copying to a key-sequenced data set, the records to be copied must be in ascending order with no duplicates (within the input data set). With an entry-sequenced data set, the records to be copied can be in any order.

REPRO causes Access Method Services to retrieve record numbers from a sequential, indexed-sequential, or VSAM data set and store them in VSAM format in key sequence, record-number sequence, or entry sequence, or store them in a sequential data set. When records are stored in key sequence, index entries are created and loaded into an index component as control intervals and control areas are filled up. Free space, as indicated in the data-set definition in the catalog, is left, and records are stored on particular volumes according to key ranges, if indicated in the definition.

You can load all of the records in one job or load them in several jobs. In subsequent jobs VSAM continues to store records as before, extending the data set as required.

To use your own program to load a key-sequenced data set, first sort the records (or build them) in key sequence, and store them with sequential access (the PUT macro). See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide to find out how to use the VSAM macros to write your own program to load records into a data set.

COPYING A CATALOG

The REPRO command can be used to copy a catalog from one volume to another. When using REPRO to copy a catalog, both the input and the output objects must be catalogs (that is, you must first define a catalog on the device that is to contain the newly-copied catalog). A catalog might be copied to move it to a faster device type or to optimize the catalog's allocation.

The receiving catalog need not be of the same device type as the source catalog. The receiving catalog must be empty. That is, the receiving catalog cannot contain any entries other than the entries that describe the catalog and its data space.

You cannot copy the contents of a nonrecoverable catalog into a recoverable catalog, and vice versa, nor can you copy a recoverable catalog into a recoverable catalog. To convert a nonrecoverable catalog into a recoverable catalog, you must export each VSAM data set (from the nonrecoverable catalog), then import it into the newly defined (recoverable) catalog. To convert a recoverable catalog into a nonrecoverable catalog or to copy a recoverable catalog into a recoverable catalog, you can either export each VSAM data set from the source catalog and then import it into the target catalog or use the EXPORTRA/IMPORTRA commands to accomplish the same thing on a volume basis.

If the specified input and output catalogs are the same catalog, the free chain rebuild in-place function is performed. This is the only function performed in this case. All DD cards referring to the catalog, including the JOBCAT/STEP CAT, must have DISP=SHR coded for this function. In addition, the rebuild function requires the INFILE(ddname)/OUTFILE(ddname) form of REPRO. The INDATASET/OUTDATASET (dynamic allocation) form of REPRO may not be used.

The free chain rebuild in-place function is allowed for both recoverable and nonrecoverable catalogs.

To use the REPRO command to copy a catalog, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

COPY-CATALOG PREPARATION

In preparation for copying a catalog, you should determine the amount of space to be allocated in the receiving catalog. Part of the process of determining how to allocate space for the receiving catalog includes reviewing how space is allocated in the catalog to be copied.

There is no method (besides dumping the catalog) that enables you to determine the size of the catalog's index set or high-key range. You can however, determine the size (that is, number of data records) of the low-key range. (When you define a catalog and include DATA(RECORDS...), the number of records you specify is the number of data records you want the low-key range to contain, unless you also specify an amount of space as a subparameter of INDEX.) One of the catalog's self-describing records (in the low-key range) is called the catalog-control record, or CCR. The CCR describes the free (or unused) control intervals in the catalog, and contains statistics about how the catalog is used. To print the CCR, issue the PRINT command (as shown below), specifying the name of the catalog as the input data set.

```
PRINT      INDATASET(catname) -  
           SKIP(3) -  
           COUNT(1)
```

When catname identifies a user catalog, your job must include a JOBCAT or STEPCAT DD statement to describe and allocate the catalog.

The CCR is printed in the "dump" format (an example of this format is included with the first example in "PRINT Examples"). The record's first 44 bytes are its key value. The next byte is the character "L," which identifies the record as a CCR-type record. The next nine bytes are three 3-byte fields whose hexadecimal values specify:

Displacement	Symbol	Contents
45 (X'2D')	A	The highest control interval number that can be assigned.
48 (X'30')	B	The next control interval number to be assigned.
51 (X'33')	C	The number of deleted records (that is, the number of control intervals that are unused because the entry they contained was deleted).

To determine the number of records in the low-key range, you also need information contained in the catalog's data-component record in addition to the information contained in the CCR. To list the data-component record, issue the LISTCAT command as shown below:

```
LISTCAT ENTRIES(catname)ALL
```

Your job must include a JOBCAT or STEPCAT DD statement to describe and allocate the catalog. The name of the catalog's data-component record is shown in the listing as VSAM.CATALOG.BASE.DATA.RECORD.

The number of records in the catalog is the same as the number of control intervals in the low-key range that currently contain either catalog records (both base and extension records) or free records. The formula to derive the number of records varies depending on whether the low-key range has only one extent or has more than one extent. Using the LISTCAT output, find the VOLUME group for the low-key range of the data component. This is the information listed under the VOLUME heading where the LOW-KEY is shown as 00 and the HIGH-KEY is shown as 3F. Extents are shown under the EXTENTS heading for this VOLUME group. By looking at the EXTENTS information you can determine whether the low-key range has one or more than one extent.

1. If the low-key range has one extent, the value of B obtained from the CCR is the number of records to use.
2. If the low-key range has more than one extent, use the formulas shown below to obtain the number of records:

- a. For all low-key range extents except the last:

$$((HIGH-RBA + 1) - LOW-RBA) / 512$$

- b. For the last extent:

$$((B \times 512) - LOW-RBA) / 512$$

where:

HIGH-RBA and LOW-RBA are the values obtained from the LISTCAT output for each extent, and B is the value obtained from the CCR.

The sum of each a plus b, above, is the number of records to use.

To the number of records derived from either 1 or 2 above, you should add the constant 19. This will allow for the fifteen self-describing records of the target catalog plus the four free records which are always reserved for VSAM catalog use at the end of the low-key range. The final result should be specified as the value for primary when you specify:

```
DATA(RECORDS(primary[b secondary]))
```

on your DEFINE USERCATALOG command.

You should keep in mind that the number of records derived from the formulas above includes any formatted free records. If the low-key range has only one extent, all control intervals which have never been used plus all control intervals that are unused because the entry they contained was deleted are formatted free records. The value contained in the CCR field described as C above includes both types of free records. If the low-key range has more than one extent, only those control intervals that are unused because the entry they contained was deleted, are formatted free records. This is the value contained in the CCR field described as C above. All formatted free records are copied from the source catalog to the target catalog by the copy catalog function.

When you specify the space allocation for the receiving catalog in terms of RECORDS as described above, you are defining the smallest possible receiving catalog. The value contained in the CCR field described as C above of the source catalog will determine how many free records will exist in the receiving catalog. These records will be available to hold new entries.

The symbols A, B, and C have no meaning except in the formulas described above, which yield the number of control intervals that contain catalog records in the low-key range. The complete format of the CCR record is included in Virtual Storage Access Method (VSAM) for OS/VS2 MVS: Catalog Management Logic. See Planning for Enhanced VSAM Under OS/VS for more details about the catalog's structure and the low-key range.

COPY-CATALOG PROCEDURE

The following steps should be followed to copy a catalog:

1. Use the DEFINE command to define a catalog into which the source catalog is to be copied. The preceding section describes a method of determining in records the smallest possible receiving catalog.

The receiving catalog need not be of the same device type as the source catalog. The receiving catalog must be empty. That is, the receiving catalog cannot contain any entries other than the entries that describe the catalog and its data space.

2. Use the REPRO command to copy the source catalog to the receiving catalog. You cannot specify any of the REPRO command's delimiters (FROMKEY, TOKEY, SKIP, or COUNT) when you use REPRO to copy a catalog. A concatenated JOBCAT or STEPCAT DD statement is required to describe and allocate both the source and receiving catalogs. When the master catalog is being copied (that is, the source is the master catalog), use a single JOBCAT or STEPCAT DD statement to describe and allocate the receiving catalog. (A DD statement is not required to describe and allocate the master (source) catalog. If a DD statement is included, an error message will result.)
3. Use the EXPORT command to disconnect the source user catalog from the master catalog. This step is not necessary if the source catalog is the master catalog.

If the source catalog was a user catalog, you should take care to list the aliases of the user catalog in the master catalog before you issue the EXPORT command. You can use the LISTCAT command to list the aliases:

```
LISTCAT -  
  CATALOG(mastercatname/password) -  
  ENTRIES(usercatname) -  
  ALL
```

The user catalog's aliases are listed under the heading "ASSOCIATIONS."

If you use the IMPORT command to make the catalog available to another system, you can issue the DEFINE ALIAS command to reestablish each of the catalog's aliases.

4. If the source catalog was a user catalog, use the DEFINE ALIAS command to establish the required aliases of the receiving catalog.
5. Use the DELETE command to remove the source catalog from the receiving volume. The source catalog appears in the receiving catalog as a cluster as a result of the copy operation.

The DELETE CLUSTER to remove the source catalog records from the receiving catalog is a special form of DELETE. It results in the following:

- The source catalog records are deleted from the target catalog.
- The source catalog's Format-1 DSCB is changed to a suballocable data space.
- The volume that contained the source catalog is now owned by the receiving (target) catalog.

After this delete, the source catalog no longer exists. If you do not want the new catalog to own the source catalog volume, use DELETE SPACE for that volume to remove volume ownership from the new catalog.

The copy-catalog procedure should be performed with caution. Until the source catalog is disconnected from the master catalog, two catalogs are available for use.

At the completion of the copy-catalog procedure:

- The DSCB on the source-catalog volume is modified so that it no longer indicates that the volume contains a catalog (step 5).
- A volume entry is created in the receiving catalog for the source-catalog volume; the source-catalog volume is owned by the receiving catalog (step 2).
- Space used by the source catalog is made available for suballocation (step 5).

If you are copying the master catalog, special care should be taken that the job is executed when the system is otherwise quiesced. Otherwise, it is possible that updates could be made to the master catalog during the copy. To use the target catalog as a master catalog, the following steps should be taken following step 2 described above:

- Replace the SYSCATLG member in SYS1.NUCLEUS (now cataloged in the target master catalog) so that the SYSCATLG member points to the new master catalog ("Appendix G. Changing the Attributes of the VSAM Master Catalog" gives an example of how to replace SYSCATLG).

- ReIPL the system. With the SYSCATLG member changed, the IPL will reference the new master catalog.
- Proceed with step 5 described above.

BACKING UP A CATALOG

You can use the REPRO command to unload a VSAM catalog into a key-sequenced, entry-sequenced, or sequential (SAM) data set. If the catalog becomes inaccessible, you can redefine the catalog and use REPRO to reload the backup to replace it. If the catalog is accessible, you can use REPRO to reload the backup and reestablish catalog entries. Reloading a catalog should be done judiciously since it is difficult to recover VSAM data spaces, page spaces, and data sets that have extended after a backup copy of the catalog was made.

Using REPRO to unload or reload a password-protected catalog requires the catalog's master password.

The parameters that limit the extent of copying are invalid for unload/reload. Parameters that indicate action on the output data set are also invalid. These parameters are: COUNT, FROMKEY, FROMNUMBER, FROMADDRESS, SKIP, TOKEY, TOADDRESS, TONUMBER, REUSE and REPLACE.

Use a STEPCAT DD statement to identify the catalog that you are loading or reloading.

Do not allow the catalog to be updated during the operation. (REPRO does not itself prevent other processing of the catalog.)

To use the REPRO command to unload or reload a catalog, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

UNLOADING A CATALOG

A sequential, key-sequenced, or entry-sequenced backup copy of a catalog is inaccessible as a catalog. Because there is little advantage in having the backup on a direct-access volume, you can most conveniently use magnetic tape to copy a catalog in a sequential data set.

To unload a catalog into a key-sequenced or entry-sequenced data set, first define the data set in another catalog (for protection).

To continually have a recent backup copy available, you should unload a catalog periodically. With tape, you can easily alternate two or more volumes for several levels of backup.

Use LISTCAT before unloading a catalog. You can compare the listing with the one you obtain after reloading. If you are unloading a recoverable catalog, use LISTCRA with the COMPARE option to ensure that the catalog and its volumes are synchronized at the time of the unload operation.

Example 3 in the section "REPRO Examples" describes the DCB parameters you must specify on your DD statement if you unload your catalog to a nonVSAM sequential (SAM) data set.

RELOADING A CATALOG

You must use REPRO to reload the backup copy into a catalog (the "target") with the same name, volume-serial number, and device type as the original catalog. Reloading the master catalog has special requirements (for details, see the section "Backing Up the Master Catalog" in the chapter "Data Security and Protection.")

The target catalog can be either a version of the original catalog (which might have been obtained using IEHDASDR RESTORE), or a newly defined user catalog (after using EXPORT DISCONNECT to remove the user catalog entry in the Master Catalog). The primary allocation of the newly defined catalog must be able to hold at least as many records as the original catalog held at the time the backup copy was made. Catalog extension will not take place during the reload operation. If the newly defined catalog is not large enough, your reload job will fail. The same method as that described in "Copy-Catalog Preparation" in this chapter can be used to determine the smallest possible size of the newly defined catalog. However, you must make the calculations based on the original catalog at the time the backup was made.

Reloading a version of the original catalog results in a catalog equivalent to the original one at the time the backup was made. Reloading replaces entries in the target catalog with entries of the same name in the backup. It inserts into the target entries that exist only in the backup. It deletes from the target entries that exist only in the target. During reloading, Access Method Services issues a maximum of one hundred messages to indicate entries that exist only in the target or only in the backup.

Reloading a newly defined catalog has the same results as reloading a version of the original catalog, with one exception. The newly defined catalog's Volume record contains only self-defining information. A check is initiated when reload opens the catalog and, if the catalog is new, a reload of the original version of the Volume record is bypassed. (The assumption is that, if the catalog is new, no other VSAM data space exists on the volume under normal conditions.) The Volume record of the original version of the catalog contains all the data space information previously on the volume at the time the catalog was unloaded. But, when the reload of the Volume record is bypassed, the data space information is lost. After the reload of a newly defined catalog, the entries in the reloaded catalog-reference data sets previously existing on the volume. If these data sets still exist, they can be accessed, but any attempt to extend the data space in which they reside will fail. In this situation, you can restore all needed information to the Volume record by, first, using EXPORT PERMANENT to remove the data set entries from the new catalog, then define a data space large enough to accommodate the data sets, and then use IMPORT to put the data sets into the newly defined data space.

Note: LISTCAT cannot run in a job step in which the catalog is empty when opened. To ensure that the LISTCAT correctly reflects the contents of the catalog, it is recommended that it be run as a separate job step.

After you reload a catalog, use LISTCAT to list its contents. Run LISTCAT in a separate job step, so that the catalog will be closed after it is reloaded (to update its self-defining information). To ensure that you have used the right backup, compare the listing with the one you obtained before unloading the original catalog.

Reloading or restoring a recoverable catalog does not cause the catalog recovery area to be updated. Therefore, the LISTCRA command with COMPARE option should be run to identify mismatches between the catalog and the catalog recovery area (CRA). (This should be done in a separate job step, immediately after the reload.) These mismatches should be resolved as necessary before

the catalog can be used. (See the section "Catalog Recovery" in the "Data Security and Protection" chapter.) No other jobs should be run after the reload and before the LISTCRA if these jobs access any of the data sets cataloged in the reloaded catalog.

If VSAM data sets or data spaces have been deleted or permanently exported since the last catalog backup and the catalog is reloaded or restored, then the deleted data sets or data spaces will still be defined in the restored catalog. Any attempt to process these entries will yield unpredictable results because the space reflected in the catalog may no longer be owned by the catalog. The catalog may be corrected by reissuing the DELETE commands.

If VSAM data sets or data spaces have been defined or imported since the last catalog backup and the catalog is reloaded or restored, then the defined data sets or data spaces will not be defined in the reloaded or restored catalog. Processing these data sets or data spaces by means of the restored catalog is not possible since they cannot be accessed. The space formerly occupied by these VSAM data sets or data spaces will not be usable, but may be recovered by scratching the format-1 DSCBs in the VTOC for the data spaces. If any volumes were added to the catalog (between the backup and the recovery), they will also be unusable until you use the DELETE command with FORCE option or ALTER REMOVEVOLUMES to give up volume ownership.

If a VSAM data set has been extended since the last catalog backup, the new extents will not be defined in the restored or reloaded catalog. Any attempt to process records in the added extents will result in a logical error. If the data set has been extended within space already allocated to the data set before the backup but has acquired no new extents, then you can issue the VERIFY command to update the catalog pointers, and the data set may be accessed normally.

The data in any extents that have been acquired by the data set since the catalog was backed up is unrecoverable. For an entry-sequenced data set the data in any new extents should consist only of records that have been added to the end of the data set. Therefore, it is possible to recover all of the data in the old extents by accessing the data set sequentially up to the end of the old physical space allocation. For a key-sequenced data set, the new extents may be any portion of the data set because of control-area splits. An attempt to read the data in logical sequence will fail with an invalid RBA indication when the data in the new extents is reached. You could access the key-sequenced data set by means of address sequence, but you then have the problem of identifying the missing records. Individual data set recovery for those data sets affected will be necessary.

See the section "Updating a Backup Catalog" in the chapter "Data Security and Protection" for a discussion of making the contents of the backup catalog agree with the contents of the original catalog at the time it became inaccessible.

OPTIMIZING THE PERFORMANCE OF UNLOAD/RELOAD

You can specify additional I/O buffers for unloading and reloading by using:

- The AMP parameter in the STEPCAT DD statement that identifies the catalog—AMP= 'BUFND=x, BUFNI=2', where x equals 2 times the number of 512-byte control intervals per track of the device used for the catalog.
- The AMP parameter in the DD statement that identifies a key-sequenced or entry-sequenced backup copy—AMP= 'BUFND=x, BUFNI=2', where x equals 2 times the number of 512-byte control intervals per track of the device used for the backup.

- The DCB parameter in the DD statement that identifies a sequential backup copy—DCB=BUFNO=x, where x equals either:
 - 2 times the number of 512-byte control intervals per track of the device used for the catalog (when the backup is on magnetic tape) or
 - 2 times the number of physical records per track of the device used for the backup (when the backup is on a direct-access volume).

Block the records in a sequential backup data set. Some catalog records are 47 bytes long; the rest are 505 bytes long. Use DCB=RECFM=VB.

PRINTING DATA SETS

You use the PRINT command to list part or all of a key-sequenced, relative-record, or entry-sequenced VSAM data set, an alternate index, a VSAM catalog, or a nonVSAM data set. The components of a key-sequenced data set or alternate index can be listed individually. To list a component of a key-sequenced data set or alternate index, specify the component name as the data set name. An alternate index is printed as though it were a key-sequenced cluster.

If a sequential or indexed-sequential data set is not cataloged, you must include the appropriate volume and unit parameters on your DD statements.

You must also supply a minimum set of DCB parameters when the data set to be printed is sequential or indexed-sequential. The following table shows the four key parameters:

Parameters	Must Be Supplied By User	Default If Not Supplied
DSORG	IS	PS
RECFM	F, FB, V, VB	U
BLKSIZE	blocksize	None
LRECL	lrecl	BLKSIZE for F or FB BLKSIZE-4 for V or VB

Note that the one parameter which is not supplied by default is BLKSIZE; you must supply this value. The DCB parameters can be supplied via the DSCB or header label of a standard labeled tape; otherwise, they must be supplied via the DD statement.

Access Method Services does not support records greater than 32,760 bytes for nonVSAM data sets (that is, LRECL=X is not supported). If the logical record length of a nonVSAM input data set is greater than 32,760 bytes, your PRINT command will terminate with an error message.

Sequential and entry-sequenced data sets are listed in physical sequential order. Indexed-sequential and key-sequenced data sets can be listed in key order or in physical sequential order. A base cluster can be listed in alternate-key sequence by specifying a path name as the data-set name for the cluster.

Only the data content of logical records is listed. System defined control fields are not listed. Each record listed is identified by one of the following:

- Its relative byte address (RBA) for entry-sequenced data sets.
- Its key for indexed-sequential (ISAM) and key-sequenced data sets, and for alternate indexes.

- Its sequential record number for sequential (nonVSAM) and relative-record data sets.

Note: If four logical and/or physical errors are encountered while trying to read the input, the printing is terminated.

To use the PRINT command to print a VSAM catalog, Access Method Services must be authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor for information about program authorization.

Refer to the chapter "Command Format" for a description of the command format for PRINT. Examples of the PRINT command are included.

LISTING TAPE VOLUMES MOUNTED AT CHECKPOINT

OS/VS Checkpoint/Restart explains taking checkpoints and restarting programs. Here is a summary:

During processing, a program can issue the CHKPT macro to record various information for use in restarting the processing if an error prevents the program from continuing. Recording information by way of CHKPT is called taking a checkpoint. The records that contain the information make up a checkpoint entry in the checkpoint data set, which contains an entry for each checkpoint that is taken.

The checkpoint data set can be a sequential data set or a partitioned data set. In a partitioned data set, each checkpoint entry is a member of it.

Checkpoint information includes the volume serial numbers of tape data sets that were open at the checkpoint. The CHKLIST command enables you to list these volume serial numbers to identify the tape data sets that need to be mounted for restart.

For a checkpoint data set with DSORG=PS (sequential data set), you can select one or more specific checkpoint entries for which the tape information is to be listed by a single CHKLIST command. By not selecting any specific entry, all checkpoint entries will be processed.

You can use CHKLIST to process a checkpoint data set with DSORG=PO (partitioned data set) in the following manner:

- Specify DSNAME=dsname(member) on the JCL statement that defines the checkpoint data set.
- Do not select a specific checkpoint entry, so that the single entry specified by member after dsname will be processed.

The CHKLIST command causes the following information to be listed:

- The checkpoint identifier for the entry being processed.
- For each tape data set that was open at the time of the checkpoint, the following items are listed:
 - dsname
 - DD name
 - type of unit on which the volume was mounted
 - the sequence number of the mounted volume
 - volume serial numbers with an * by the volume serial number of the mounted volume.

To process multiple members of a partitioned checkpoint data set, use the CHKLIST command once for each member.

Note: The CHKLIST command cannot be invoked as a TSO command.

Refer to the chapter "Command Format" for a description of the command format for CHKLIST. Examples of the CHKLIST command are also included.

COMMAND FORMAT

This chapter sets out the functional command formats. Parameters are grouped into categories such as name, data organization, allocation, and protection and integrity. The format of each command is then shown, followed by a discussion of each parameter. Examples of each command follow the parameter discussion.

FUNCTIONAL COMMAND FORMAT

This section provides complete reference information about all functional commands of Access Method Services. The commands discussed in the section are:

- ALTER command, which is used to alter attributes of data sets and other objects that have already been defined.
- BLDINDEX command, which is used to build an alternate index over a base cluster.
- CHKLIST command, which lists tape data sets opened during a checkpoint.
- CNVTCAT command, which converts OS catalog entries to VSAM catalog entries and merges them into a VSAM catalog.
- DEFINE ALIAS command, which is used to define an alternate name for a nonVSAM data set or a user catalog.
- DEFINE ALTERNATEINDEX command, which is used to define an alternate index.
- DEFINE CLUSTER command, which is used to define a cluster for a key-sequenced, entry-sequenced, or relative-record data sets.
- DEFINE GENERATIONDATAGROUP command, which is used to create a catalog entry for a generation data group.
- DEFINE NONVSAM command, which is used to define a catalog entry for a nonVSAM data set.
- DEFINE PAGESPACE command, which is used to define an entry for a page space data set.
- DEFINE PATH command, which is used to define a path directly over a base cluster or a path over an alternate index and its related base cluster.
- DEFINE SPACE command, which is used to define a VSAM data space.
- DEFINE USERCATALOG command, which is used to define a VSAM user catalog.
- DELETE command, which is used to delete data sets and other objects, including catalogs and nonVSAM data sets.
- EXPORT command, which is used to export VSAM files and to disconnect user catalogs.
- EXPORTRA command, which is used to recover VSAM and nonVSAM catalog entries from catalog recovery areas and, for VSAM objects (clusters and alternate indexes), to recover the data itself by means of catalog recovery areas.

- IMPORT command, which is used to import VSAM data sets and to connect user catalogs.
- IMPORTRA command, which is used to reconstruct multiple VSAM data sets from a data set created by the EXPORTRA command.
- LISTCAT command, which is used to list catalog entries.
- LISTCRA command, which is used to list or compare the contents of a given catalog recovery area.
- PRINT command, which is used to print both VSAM and nonVSAM data sets.
- REPRO command, which is used to copy and load both VSAM and nonVSAM files and VSAM catalogs.
- VERIFY command, which is used to verify and correct certain problems that have made your data set unusable.

See "Notation Conventions" in the "Preface" (About This Book) for an explanation of the symbols used in the command formats.

ALTER

The format of the ALTER command is:

ALTER	<pre>entryname[/password] [ADDVOLUMES(volser[b volser...])] [ATTEMPTS(number)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLPW(password)] [DESTAGWAIT NODESTAGWAIT] [EMPTY NOEMPTY] [ERASE NOERASE] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE(CI-percent[b CA-percent])] [INHIBIT UNINHIBIT] [KEYS(length b offset)] [MASTERPW(password)] [NEWNAME(newname)] [NULLIFY([AUTHORIZATION(MODULE STRING)] [CODE] [CONTROLPW] [EXCEPTIONEXIT] [MASTERPW] [OWNER] [READPW] [RETENTION] [UPDATEPW])] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE(average b maximum)] [REMOVEVOLUMES(volser[b volser...])] [SCRATCH NOSCRATCH] [SHAREOPTIONS(crossregion[b crosssystem])] [STAGE BIND CYLINDERFAULT] [TO(date) FOR(days)] [UNIQUEKEY NONUNIQUEKEY] [UPDATE NOUPDATE] [UPDATEPW(password)] [UPGRADE NOUPGRADE] [WRITECHECK NOWRITECHECK] [CATALOG(catname[/password])]</pre>
-------	---

ALTER Parameters: Summary

The parameters of the ALTER command are described in the following groups:

- Name, which describes the entryname and NEWNAME parameters. These parameters are used to name the entry to be altered and, optionally, to rename the entry. These parameters are also used to rename a member of a nonVSAM partitioned data set.
- Protection and integrity, which describes the NULLIFY, MASTERPW, CONTROLPW, UPDATEPW, READPW, UNINHIBIT, INHIBIT, CODE, ATTEMPTS, AUTHORIZATION, EXCEPTIONEXIT, OWNER, TO, FOR, SHAREOPTIONS, ERASE, NOERASE, ERASE, NOERASE, WRITECHECK, NOWRITECHECK, DESTAGWAIT, and NODESTAGWAIT parameters. These parameters are used to alter protection and integrity attributes.
- Allocation, which describes the FILE, FREESPACE, BUFFERSPACE, RECORDSIZE, KEYS, ADDVOLUMES, REMOVEVOLUMES, STAGE, BIND, and CYLINDERFAULT parameters. These parameters are used to modify the amount of free space to be left in

control intervals and control areas, to modify the amount of buffer space to be provided, to add and remove volumes from the list of volumes to be used as overflow volumes, and to modify the indication of how a data or index component that is stored on a mass-storage volume is to be staged.

When a data set does not contain any data records, RECORDSIZE can be specified to modify the data set's average and maximum record sizes. When a key-sequenced data set or alternate index is empty, KEYS can respecify the the length and position of each data-record's key field.

- Alternate index and path, which describes the UNIQUEKEY, NONUNIQUEKEY, UPDATE, NOUPDATE, UPGRADE, and NOUPGRADE parameters. These parameters specify attributes that apply only to alternate index and path entries.
- Generation-data-group attributes, which describes the EMPTY, NOEMPTY, SCRATCH, and NOSCRATCH parameters. These parameters are used to alter what is to happen when the maximum number of generation data sets is reached.
- Catalog, which describes the CATALOG parameter. This parameter is used to name the catalog in which the entry to be altered is defined.

Entry-Types That Can Be Altered

An "X" in Figure 8 indicates you can respecify the value or attribute for that type of VSAM catalog entry. Some attributes, when specified for a cluster or alternate index, apply to its data or index component entry only, and not to the cluster or alternate-index entry. Some attributes can only be specified for the cluster's or alternate-index's data or index component entry—you must identify the component's entry with its entryname. You can use the Access Method Services LISTCAT command to determine the names VSAM generated for the object's components.

When you identify a group of entries with a generic name, entries whose entrynames match the supplied qualifiers are altered, if they contain the type of information specified with the ALTER command's other parameters.

ADDVOLUMES		X	X		X	X						
ATTEMPTS	X	X	X	X	X	X	X	X	X			
AUTHORIZATION	X	X	X	X	X	X	X	X	X			
BIND		X	X		X	X						
BUFFERSPACE		X			X				X			
CODE	X	X	X	X	X	X	X	X	X			
CONTROLPW	X	X	X	X	X	X	X	X	X			
CYLINDERFAULT		X	X		X	X						
DESTAGEWAIT		X	X		X	X						
EMPTY												X
ERASE		X			X							
EXCEPTIONEXIT		X	X		X	X						
FILE	X	X	X	X	X	X	X	X	X	X		
FOR	X			X			X	X	X	X		X
FREESPACE		X			X							
INHIBIT		X	X		X	X						
KEYS	X	X		X	X			X				
MASTERPW	X	X	X	X	X	X	X	X	X			
NEWNAME	X	X	X	X	X	X	X	X		X		
NODESTAGEWAIT		X	X		X	X						
NOEMPTY												X
NOERASE		X			X							
NONUNIQUEKEY		X										
NOSCRATCH												X
NOUPDATE								X				
NOUPGRADE	X											
NOWRITECHECK		X	X		X	X						
NULLIFY	X	X	X	X	X	X	X	X	X	X	X	X
AUTHORIZATION	X	X	X	X	X	X	X	X	X			
CODE	X	X	X	X	X	X	X	X	X			

Figure 8 (Part 1 of 2). ALTER Parameters and the Entry-Types to Which Each Applies

Type of Catalog Entry:

	ALTERNATE INDEX		CLUSTER		DATA INDEX		PGSPC	PATH	UCAT	NonVSAM	GDG
	DATA	INDEX	DATA	INDEX	DATA	INDEX					
CONTROLPW	X	X	X	X	X	X	X	X	X		
EXCEPTIONEXIT		X	X		X	X					
MASTERPW	X	X	X	X	X	X	X	X	X		
MODULE	X	X	X	X	X	X	X	X	X		
OWNER	X	X	X	X	X	X		X	X		
READPW	X	X	X	X	X	X	X	X	X		
RETENTION	X			X			X	X	X	X	X
STRING	X	X	X	X	X	X	X	X	X		
UPDATEPW	X	X	X	X	X	X	X	X	X		
OWNER	X	X	X	X	X	X		X	X		
READPW	X	X	X	X	X	X	X	X	X		
RECORDSIZE	X	X		X	X			X			
REMOVEVOLUMES		X	X		X	X					
SCRATCH											X
SHAREOPTIONS		X	X		X	X					
STAGE		X	X		X	X					
TO	X			X			X	X	X	X	X
UNINHIBIT		X	X	X	X						
UNIQUEKEY		X									
UPDATE								X			
UPDATEPW	X	X	X	X	X	X	X	X	X		
UPGRADE	X										
WRITECHECK		X	X		X	X	X				

Figure 8 (Part 2 of 2). ALTER Parameters and the Entry-Types to Which Each Applies

You cannot identify the following types of entries: alias entries, a catalog's data or index-component entries, and a master-catalog's self-describing entries.

ALTER PARAMETERS

Required Parameters

entryname[/password]
is a required parameter that names the entry to be altered and supplies a password. See the section "Generic Names and ALTER" for a description of how to specify a generic name in order to alter multiple entries with one ALTER command.

See "Modifying Catalog Information" if you wish to use data or index entryname or cluster entry subparameters.

If you are renaming a member of a nonVSAM partitioned data set, the entryname must be specified in the format: pdsname(membername).

If you are altering a password-protected entry in a password-protected catalog, you must specify a password. The password can be specified with entryname or in the CATALOG parameter. For the following entry types, the password must be the master password for the entry or for the catalog that contains the entry: alternate index, cluster, page space, path, user catalog, data, and index. If a data or index component entry is to be altered, the master password of the cluster, component, or catalog can be supplied. For the following entry types, the password must be the update- or higher-level of the catalog that contains the entry: nonVSAM and GDG.

Optional Parameters

ADDVOLUMES(volser[b volser])
specifies volumes to be added to the list of overflow (candidate) volumes.

volser
A volume serial number, volser, may contain 1 to 6 alphanumeric, national (@, #, and \$), and special characters; the special characters include commas, blanks, semicolons, parentheses, slashes, asterisks, periods, quotation marks, ampersands, plus signs, hyphens, and equal signs. A volume serial number must be enclosed in single quotation marks if it contains a special character. Single quotation marks within a volume serial number must be coded as two single quotation marks.

The volumes to be added as candidate volumes must already be owned by the catalog that contains the entry being altered; that is, space must have been defined on a volume to be added or the volume must have been identified as a candidate volume.

Abbreviation: AVOL

ATTEMPTS(number)
specifies the maximum number of times the operator can try to enter a correct password in response to a prompting message. If 0 is specified, the operator is not prompted and is not allowed to enter a password from the console. This parameter can be coded, but only has effect when the entry's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

number
is an integer from 0 to 7 and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Note to TSO users: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password. Checking the logon password counts as one attempt to obtain a password. If ATTEMPTS is not specified, the user has one attempt to supply the password because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])

specifies that a user-security-verification routine (USVR) is available for additional security verification.

entrypoint

specifies the name of the USVR. entrypoint can contain 1 to 8 alphanumeric characters (A through Z, and 0 through 9), or national characters (@, #, \$), and special characters (the hyphen (-) and the 12-0 overpunch (X'C')). The name's first character is either an alphabetic character or a national character.

string

specifies information to be passed on to the USVR when it receives control to verify authorization. string can contain 1 to 255 EBCDIC characters.

When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], semicolon [;], ampersand [&], plus [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single quotation marks (for example, AUTHORIZATION(entrypoint, '*DORIS*')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, '*CA' 'ROL*')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint, SUE). The string can contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

BUFFERSPACE(size)

specifies the minimum space to be provided for buffers. The amount specified should be greater than or equal to the amount specified in the original definition. If the amount is less than was specified when the entry was defined, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key sequenced, one index component control interval. BUFFERSPACE can be specified only for a catalog, or for a cluster's or alternate-index's data component.

size

is the amount of space to be provided for buffers. If you are altering the bufferspace for a catalog, the decimal values you can specify are limited to 3072, 4096, 5120, 6144, 7168, and 8192.

size can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. If the size specified is less than the amount VSAM requires, VSAM gets the amount it requires when the data set is opened.

Abbreviations: BUFSP or BUFSPC

CATALOG(catname[/password])

specifies the catalog location of the entry to be altered. See "Order of Catalog Use: ALTER" for information about the order in which catalogs are searched.

catname

specifies the name of the catalog that contains the entry.

password

specifies the master password of the catalog that contains the entry to be altered. If the entry to be altered is password protected and the catalog is also password protected, a password must be entered either through this parameter or through the parameter that specifies the entry to be altered.

Abbreviation: CAT

CODE(code)

specifies a code name for the entry being altered. If an attempt is made to access a password-protected entry without a password, the code name is used in a prompting message; the code enables the operator or TSO terminal user to be prompted for the password without disclosing the name of the entry.

If CODE is not specified and an attempt is made to access a cluster or component that is password protected without supplying a password, the operator or TSO terminal user is prompted with the name of the entry.

code

can contain 1 to 8 alphameric or special characters. When the code contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*], enclose the code in single-quotation marks (for example, CODE('*DORIS*')).

If the code contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE('*CA''ROL*')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CODE(X'E2E4C5') is the same as CODE(SUE).

CONTROLPW(password)

specifies a control password for the entry being altered.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*], enclose the password in single-quotation marks (for example, CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('X'CA'ROLX')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

DESTAGEWAIT|NODESTAGEWAIT

specifies whether a data or index component that is stored on a mass-storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it. These parameters can be specified only for the data or index component of a cluster or alternate index.

DESTAGEWAIT indicates that destaging is to be completed before control is returned from VSAM to the program that issues the CLOSE macro. VSAM can thus notify the program whether destaging was successful.

NODESTAGEWAIT indicates that notification of unsuccessful destaging is to be made only by a message to the operator.

Abbreviations: DSTGW and NDSTGW

EMPTY|NOEMPTY

specifies what is to happen when the maximum number of generation data sets has been cataloged. EMPTY specifies that all of the generation data sets are to be uncataloged. NOEMPTY specifies that only the oldest generation data set is to be uncataloged.

Abbreviations: EMP and NEMP

ERASE|NOERASE

specifies whether the data component is to be erased when its entry in the catalog is deleted. If ERASE is specified, the component is overwritten with binary zeros when its catalog entry is deleted.

Abbreviations: ERAS and NERAS

EXCEPTIONEXIT(entrypoint)

specifies the name of the user-written routine that receives control when an exception (usually an I/O error) occurs while the entry's object is being processed. An exception is any condition that causes a SYNAD exit to be taken. The object's exception exit routine is processed first, then the user's SYNAD-exit routine receives control.

You cannot specify this parameter to add an exception-exit routine to a VSAM cluster that is cataloged in a VSAM catalog created in an OS/VS2-2 (or lower release level) system.

Abbreviation: EEXT

FILE(dname)

specifies the name of a DD statement that describes the volume whose catalog recovery area contains the entry's copy. You can specify FILE when the entry to be altered is in a recoverable catalog. See the section "Restoring Catalog Entries After System Failure" for a description of catalog recovery area contents for each volume. The volume whose catalog recovery area contains a copy of the catalog entry must always be identified in the DD statement referenced via FILE. In addition:

- If you specify ADDVOLUMES or REMOVEVOLUMES, the volume(s) being added or removed must be identified.
- If you specify NEWNAME for a nonVSAM data set or the data or index component of a unique cluster, alternate index, or page space which does not reside on the same volume as the catalog recovery area entry, that volume(s) must be identified. If FILE is not specified, the catalog recovery area volume(s) is dynamically allocated. It must be mounted as permanently resident or reserved.

FILE can be used to specify the name of a DD statement that identifies the volume of an entry to be renamed if a nonVSAM data set or the data or index component of a unique cluster, alternate index, or page space is to be renamed. If FILE is not specified, the object's volume is dynamically allocated. The object's volume must be mounted as permanently resident or reserved.

FILE can be used to identify the name of the partitioned data set if a member is to be renamed. If FILE is not specified, the partitioned data set is dynamically allocated.

If multiple volumes of different device types are to be identified via FILE, you must use concatenated DD statements.

FREESPACE(CI-percent[b CA-percent])

specifies the amount of space that is to be left free after any allocation and after any split of control intervals (CI-percent) and control areas (CA-percent).

The amounts are specified as percentages. The percentages, which must be equal to or less than 100, may be expressed in decimal, hexadecimal, or binary. If you specify 100 percent of free space, one record is placed in each control interval and one control interval is placed in each control area.

CI-percent and CA-percent can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviation: FSPC

INHIBIT|UNINHIBIT

specifies whether the entry being altered can be accessed for any operation or for only read operations. INHIBIT specifies that the entry being altered is only to be read. UNINHIBIT specifies that the read-only restriction set by a previous ALTER or EXPORT command is to be removed.

Abbreviations: INH and UNINH

KEYS(length b offset)

specifies the length and offset of the object's key. If the entry being altered defines an alternate index, offset applies to the alternate key in the data records in the base cluster. For more information, see entryname.

Restrictions: Can be specified only if all of the following are true:

- The object whose entry is being altered contains no data records.
- The values for KEYS in the catalog must be default values. However, if nondefault keys are altered and the new value matches the old, processing continues for any other parameters specified on the command. For default values, see the DEFINE command that defines the object.

- The key must fit within the record whose length is specified by the RECORDSIZE parameter.
- The key must fit in the first record segment of a spanned record.
- The new values for KEYS must not conflict with the control interval size specified when the object was defined.
- The object whose entry is being altered is an alternate index, a path, a key-sequenced cluster, or a data component of a key-sequenced cluster or alternate index.

length b offset

specifies the length of the key (between 1 and 255), in bytes, and its displacement from the beginning of the data record, in bytes. You can express length and offset in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

MASTERPW(password)

specifies a master password for the entry being altered. For more details about the object's master password, see the DEFINE command that defines the object.

The AUTHORIZATION, CODE, and ATTEMPTS parameters have no effect unless the entry has a master password associated with it. If MASTERPW is not specified, the highest level password specified becomes the password for all higher levels. The MASTERPW is not specified, the highest level password specified becomes the password for all higher levels. The master password allows all operations.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation: MRPW**NEWNAME(newname)**

specifies that the entry to be altered is to be given a new name.

The new name may contain 1 to 44 alphanumeric characters, (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and 12-0 overpunch). Names that contain more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of any name or name segment must be either an alphabetic character or a national character.

If the object being renamed is a unique cluster or component, or a nonVSAM data set, the object's volume must be mounted because the volume's VTOC is modified. You can use a JCL DD statement (see the FILE parameter) to cause the object's volume(s) to be allocated. If you do not supply a DD statement, the object's volume(s) is dynamically allocated. The volume(s) must be mounted as either permanently resident or reserved.

If you specify generic names, you specify both the entryname and the newname as generic names. See "Generic Names and ALTER" for details on generic names.

If you are renaming a member of a nonVSAM partitioned data set, the newname must be specified in the format: pdsname(membername).

If you are renaming a VSAM data set which is RACF-protected, the existing RACF data set profile will be renamed. If a data set profile already exists for the new data set name prior to the ALTER command, the command is terminated and the data set name and/or protection attributes remain unchanged. If the old profile cannot be found or cannot be altered to the NEWNAME, the NEWNAME action will not be completed in the catalog, and an error message will indicate the reason for noncompletion.

Abbreviation: NEWNM

NULLIFY

specifies that the protection attributes identified by subparameters of NULLIFY are to be nullified. Attributes are nullified before any respecification of attributes is performed. If all levels of passwords are nullified and none are respecified, CODE, AUTHORIZATION, and ATTEMPTS have no effect.

Abbreviation: NULL

AUTHORIZATION(MODULE|STRING)

specifies that either the user-authorization routine or the user-authorization record is to be nullified. When MODULE is specified, the module name is removed from the catalog record, but the module itself is not deleted. If you nullify the user-authorization module, the user-authorization record (character string) is also nullified. When you nullify the authorization record, the corresponding module is not nullified.

Abbreviations: AUTH, MDLE, and STRG

CODE

specifies that the code name used for prompting is to be nullified.

CONTROLPW

specifies that the control password is to be nullified.

Abbreviation: CTLPW

EXCEPTIONEXIT

specifies that the entry's exception exit is to be nullified. The module name is removed from the catalog record, but the exception-exit routine itself is not deleted.

Abbreviation: EEXT

MASTERPW

specifies that the master password is to be nullified. If a new master password is not specified and if other passwords exist, the highest level password that exists automatically becomes the password for all

higher levels, including the master.

Abbreviation: MRPW

OWNER

specifies that the owner identification is to be nullified.

READPW

specifies that the read password is to be nullified.

Abbreviation: RDPW

RETENTION

specifies that the retention period, specified in a TO or FOR parameter, is to be nullified.

Abbreviation: RETN

UPDATEPW

specifies that the update password is to be nullified.

Abbreviation: UPDPW

OWNER(ownerid)

specifies the identification of the owner of the entry being altered.

ownerid

can contain 1 to 8 EBCDIC characters. When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('*CA''ROL*')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

READPW(password)

specifies a read password for the entry being altered.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

RECORDSIZE(average b maximum)

specifies the new average and maximum lengths for the data records in the object. If the object whose entry is being altered is an alternate index path, the alternate index itself is altered; if the path points directly to its base cluster, then the base cluster is altered. If the object whose entry is being altered is an alternate index, the alternate key must be within the limit specified by maximum. For more information, see *entryname*.

Restrictions: RECORDSIZE can be specified only if all of the following are true:

- The object whose entry is being altered contains no data records.
- The value for maximum RECORDSIZE in the catalog must be the default value. However, if a nondefault value is altered and the new value matches the old, processing continues for any other parameters specified on the command. For default values, see the DEFINE command that defines the object.
- The new maximum record length must be at least seven bytes less than control interval size, unless the record is a spanned record.
- The new record length must be large enough to contain all prime and alternate keys previously defined.
- For an alternate index, if NONUNIQUEKEY is specified, RECORDSIZE must account for the increased record size resulting from the multiple prime-key pointers in the alternate index data record.
- The object whose entry is being altered is an alternate index, a cluster, a path, or a data component.

Abbreviation: RECSZ

REMOVEVOLUMES(volser[b volser])

has two uses: (1) REMOVEVOLUMES specifies volume(s) to be removed from the list of candidate volumes associated with the entry being altered. Volumes specified are removed after any new volumes are added to the candidate list. If a volume to be removed contains data that belongs to the entry being altered, the volume is not removed. For this use of REMOVEVOLUMES, the name of the data or index component must be specified in the *entryname* parameter. (2) REMOVEVOLUMES specifies volume(s) from which all VSAM data spaces are to be removed and VSAM ownership is to be taken away—without access to the user catalog that owns the volume(s). For this use of REMOVEVOLUMES, the name of the master catalog and its master password (if any) must be specified in the *entryname* parameter, and the FILE parameter is required. (You can have only one DD statement with this use of ALTER REMOVEVOLUMES—only volumes of the same device type can be processed.) See the section "VSAM Volume Cleanup" in the chapter "Data Security and Protection" and the section "VSAM Volume Recovery Function" in the chapter "Modifying Catalog Information" for information and cautions about this use of ALTER REMOVEVOLUMES.

A volume serial number, *volser*, may contain one to six alphameric, national (@, #, and \$), and special characters; the special characters include commas, blanks, semicolons, parentheses, slashes, asterisks, periods, quotation marks,

ampersands, plus signs, hyphens, and equal signs. A volume serial number must be enclosed in single quotation marks if it contains a special character. Single quotation marks within a volume serial number must be coded as two single quotation marks.

Abbreviation: RVOL

SCRATCH|NOSCRATCH

specifies whether generation data sets are to be removed from the VTOC of the volume on which they reside when they are uncataloged (that is, whether the data set's format-1 DSCB is removed—scratched—from the VTOC so that the data set can no longer be accessed).

Abbreviations: SCR and NSCR

SHAREOPTIONS(crossregion[b crosssystem])

specifies how a data or index component of a cluster or alternate index can be shared among users (see Virtual Storage Access Method (VSAM) Programmer's Guide for a full description of data set sharing). This parameter can be specified only for the data or index components of clusters and alternate indexes.

crossregion

specifies the amount of sharing allowed among regions (within the same system). Independent job steps in an operating system can access a VSAM data set concurrently. To share a data set, each user must specify DISP=SHR in the data set's DD statement. The values that can be specified are:

- 1 specifies that the data set can be shared by any number of users for read processing, or the data set can be accessed by only one user for read and write processing. With this option, VSAM ensures complete data integrity for the data set.
- 2 specifies that the data set can be accessed by any number of users for read processing and it can also be accessed by one user for write processing. With this option, VSAM ensures write integrity by obtaining exclusive control for a control interval when it is to be updated. If a user desires read integrity, it is his responsibility to use the ENQ and DEQ macros appropriately to provide read integrity for the data his program obtains.
- 3 specifies that the data set can be fully shared by any number of users. With this option, each user is responsible for maintaining both read and write integrity for the data his program accesses. User programs that ignore the write integrity guidelines can cause VSAM program checks, lost or inaccessible records, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared by any number of users and buffers used for direct processing are refreshed for each request. This option requires your program to use the ENQ and DEQ macros to maintain data integrity while sharing the data set. Improper use of the ENQ macro can cause problems similar to those described under SHAREOPTION 3.

crosssystem

specifies the amount of sharing allowed among systems. Job steps of two or more OS/VS operating systems can gain access to the same VSAM data set regardless of the disposition specified in each step's DD statement for the data set. To get exclusive control of the data set's volume, a task in one system issues the RESERVE macro. The level of cross-system sharing allowed by VSAM applies only in a multiple operating system environment. The values that can be specified are:

- 1 Reserved
- 2 Reserved
- 3 specifies that the data set can be fully shared. With this option, each user is responsible for maintaining both read and write integrity for the data his program accesses. User programs that ignore write-integrity guidelines can cause VSAM program checks, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared. Buffers used for direct processing are refreshed for each request. This option requires that you use the RESERVE and RELEASE macros to maintain data integrity while sharing the data set. Writing is limited to PUT-update and PUT-insert processing that does not change the high-used RBA if your program opens the data set with DISP=SHR. Data set integrity cannot be maintained unless all jobs accessing the data set in a cross-system environment specify DISP=SHR. Improper use of the RESERVE macro can cause problems similar to those described under SHAREOPTION 3.

Abbreviation: SHR

STAGE|BIND|CYLINDERFAULT

specifies how a data or index component that is stored on a mass-storage volume is to be staged. These parameters can be specified only for the data or index component of a cluster or alternate index.

STAGE

indicates that the component is to be staged from mass storage to a direct access storage staging drive when the component is opened.

BIND

indicates that the component is not only to be staged, but also to be bound—that is, retained on the direct access storage staging drive until it is closed.

CYLINDERFAULT

indicates that the component is not to be staged when it is opened, but that data from it is to be staged as a processing program needs it.

Abbreviation: CYLF

TO(date)|FOR(days)

specifies the retention period for the entry being altered. This cannot be specified for the data or index components of clusters, alternate indexes, or catalogs.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the number (001 through 365) of the day, through which the entry is to be kept.

FOR(days)

specifies the number of days for which the entry is to be kept. The maximum number that can be specified is 9999. If the number specified is 0 through 1830, the cluster is retained for the number of days specified; if the number is between 1831 and 9999, the cluster is retained through the year 1999.

days

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

UNIQUEKEY|NONUNIQUEKEY

specifies whether the alternate-key value can be found in more than one of the base cluster's data records. **UNIQUEKEY** specifies that each alternate-key value is unique. If the same alternate-key value is found in more than one of the base cluster's data records, an error results.

NONUNIQUEKEY allows you to have an alternate-index record that, for a given alternate-key value, might point to more than one data record in the cluster.

UNIQUEKEY can be specified only for an empty alternate index (that is, an alternate index that is defined but not yet built).

NONUNIQUEKEY can be specified for an alternate index at any time. If the alternate index is empty, you should also consider specifying **RECORDSIZE** to ensure that each alternate index record is large enough to contain more than one data record pointer.

Abbreviations: UNQK and NUNQK

UPDATE|NOUPDATE

specifies whether a base-cluster's alternate-index upgrade set is to be allocated when the path's name is allocated with a DD statement.

NOUPDATE specifies that the path's cluster is to be allocated, but that the cluster's alternate-index upgrade set is not to be allocated.

Abbreviations: UPD and NUPD

UPDATEPW(password)

specifies an update password for the entry being altered.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, **UPDATEPW('*DORIS*')**).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, **UPDATEPW('*CA''ROL*')**).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, UPDATEPW(X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

UPGRADE|NOUPGRADE

specifies whether an alternate index is to be upgraded (that is, kept up to date) when its base cluster is modified. UPGRADE specifies that when the cluster's records are added to, updated, or erased, the cluster's alternate index is upgraded to reflect the changed data.

If UPGRADE is specified when the cluster is open, the upgrade attribute does not apply to the alternate index until the cluster is closed and next opened (that is, a new set of VSAM control blocks describes the cluster and its attributes).

UPGRADE can be specified only for an empty alternate index (that is, an alternate index that is defined but not built). However, the UPGRADE attribute is not effective for the alternate index until the alternate index is built (see the BLDINDEX command).

NOUPGRADE can be specified for an alternate index at any time.

Abbreviations: UPG and NUPG

WRITECHECK|NOWRITECHECK

specifies whether a data or index component is to be checked by a machine action called write-check when a record is written into it. If WRITECHECK is specified, a record is written and then read, without data transfer, to test for the data-check condition.

Abbreviations: WCK and NWCK

ALTER EXAMPLES

Alter a Cluster's Entry: Example 1

In this example, an ALTER command is used to specify passwords for a nonindexed (entry-sequenced) cluster, D50.EXAMPLE.ESDS1. No password for the cluster is required, because the cluster was defined without passwords.

```
//ALTER1  JOB    ...
//STEP1   EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SYSIN   DD    *
          ALTER -
          D50.EXAMPLE.ESDS1 -
          MASTERPW(DEPT26M) -
          CONTROLPW(DEPT26C) -
          UPDATEPW(DEPT26U) -
          READPW(DEPT26R) -
          AUTHORIZATION(D26AUTH)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The ALTER command adds passwords to the entry-sequenced cluster's cluster catalog entry. The passwords are not added to the cluster's data entry, however. If a user's program supplies the cluster's data entry entryname and opens the data records, the unauthorized user can access the cluster's data records even though the cluster itself is password protected. The ALTER command's parameters are:

- D50.EXAMPLE.ESDS1, the name of the entry-sequenced cluster. It is assumed that an alias entry exists named D50 for the user catalog D27UCAT2. The data set name, D50.EXAMPLE.ESDS1, causes the ALTER request to be directed to D27UCAT2.
- MASTERPW, CONTROLPW, UPDATEPW, READPW, and AUTHORIZATION, which specify passwords and the entryname of the user's-security-verification routine.

Alter the Entrynames of Generically Named Clusters: Example 2

In this example, several clusters with similar names, GENERIC.*.BAKER (where "*" is any 1 to 8 character simple name), are renamed so that their entrynames are GENERIC.*.ABLE. The name "GENERIC.*.BAKER" is called a generic name.

```
//ALTER2  JOB    ...
//JOB CAT DD  DSNAME=D27UCAT1,DISP=SHR
//STEP1   EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SYSIN   DD    *
          ALTER -
          GENERIC.*.BAKER -
          NEWNAME(GENERIC.*.ABLE)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The ALTER command changes each generic entryname, GENERIC.*.BAKER, to GENERIC.*.ABLE. Its parameters are:

- GENERIC.*.BAKER, which identifies the objects to be modified.
- NEWNAME, which specifies that each generic entryname GENERIC.*.BAKER is changed to GENERIC.*.ABLE.

Alter the Attributes of a Generation Data Group: Example 3

In this example, the attributes of a generation data group are modified. Because the attributes of the group are cataloged in the generation data group's base catalog entry, only this entry is modified.

```
//ALTER3 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
ALTER -
      GDG01 -
      NOEMPTY -
      SCRATCH
/*
```

The job control language statements are:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The ALTER command modifies some of the attributes of generation data group GDG01. Its parameters are:

- GDG01, which identifies the object to be modified.
- NOEMPTY, which specifies that only the oldest generation data set is to be uncataloged when the maximum number of cataloged generation data sets is exceeded.
- SCRATCH, which specifies that the generation data set's DSCB is to be removed from the volume's VTOC when the data set is uncataloged.

The attributes specified for the generation data group with the ALTER command override any attributes previously specified for the GDG.

BLDINDEX

The format of the BLDINDEX command is:

BLDINDEX	<pre>{INFILE(dname[/password]) INDATASET(entryname[/password])} {OUTFILE(dname[/password] [b dname[/password]...]) OUTDATASET(entryname[/password] [b entryname[/password]...])} [CATALOG(catname[/password])] [EXTERNALSORT INTERNALSORT] [WORKFILES(dname b dname)]</pre>
-----------------	---

BLDINDEX can be abbreviated: BIX

BLDINDEX PARAMETERS**Required Parameters****INFILE(dname[/password])**

names the DD statement that identifies the base cluster or a path that points to the base cluster. If the base cluster or path is password protected, you can supply the read (or higher-level) password of the object named in the DD statement. Otherwise, the operator or TSO terminal user is prompted to supply the correct password.

The base cluster must be defined in the same catalog as the alternate index, and must contain data records.

The alternate index identified with OUTDATASET or OUTFILE must be related to the base cluster.

Abbreviation: IFILE

INDATASET(entryname[/password])

names a base cluster or a path that points to the base cluster. If the base cluster or path is password protected, you supply the read (or higher-level) password of the named object. Otherwise, the operator or TSO terminal user is prompted to supply the correct password. When you specify INDATASET, the base-cluster's volume is dynamically allocated. Therefore, the base-cluster's volume must be mounted as permanently resident or reserved.

The base cluster must be defined in the same catalog as the alternate index, and must contain at least one data record.

The alternate index identified with OUTDATASET or OUTFILE must be related to the base cluster.

Abbreviation: IDS

OUTFILE(dname[/password][b dname[/password]...])

names the DD statement that identifies the alternate index or a path that points to the alternate index. You can build more than one alternate index for the same base cluster by specifying more than one dname with the OUTFILE parameter. If the alternate index or path is password protected, you can supply the update- (or higher-level) password of the object named in the DD statement. Otherwise, the operator or TSO terminal user is prompted to supply the correct password.

The alternate index must be defined in the same catalog as the base cluster, and must be empty (that is, its high-used RBA equals zero) or defined with the reusable attribute.

The alternate index must be related to the base cluster identified with INDATASET or INFILE.

Abbreviation: OFILE

OUTDATASET(entryname[/password]
[b entryname[/password]...])

names the alternate index or a path that points to the alternate index. You can build more than one alternate index for the same base cluster by specifying more than one entryname with the OUTDATASET parameter. If the alternate index or path is password protected, you supply the update (or higher-level) password of the named object. Otherwise, the operator or TSO terminal user is prompted to supply the correct password. When you specify OUTDATASET, the alternate-index's volume is dynamically allocated. Therefore, the volume must be mounted as permanently resident or reserved.

The alternate index must be defined in the same catalog as the base cluster, and must be empty (that is, its high-used RBA equals zero) or must have been defined with the REUSE attribute.

The alternate index must be related to the base cluster identified with INDATASET or INFILE.

Abbreviation: ODS

Optional Parameters

CATALOG(catname[/password])

names the catalog that the workfiles are to be defined in. The workfiles are defined and used by the BLDINDEX routine. When all alternate indexes are built and the workfiles are no longer needed by the BLDINDEX routine, they are deleted.

password

If the catalog is password protected, you can supply its update- (or higher-level) password. Otherwise, the operator is prompted to supply the correct password.

See "Order of Catalog Use: BLDINDEX" for information about the order in which a catalog is selected when the CATALOG parameter is not specified.

Abbreviation: CAT

EXTERNALSORT|INTERNALSORT

specifies whether the key-pointer pairs are to be sorted entirely within virtual storage.

When you specify EXTERNALSORT, two external-sort workfiles are defined and built as VSAM entry-sequenced clusters.

When you specify INTERNALSORT, or allow it to default, Access Method Services sorts the key-pointer pairs entirely within the user-provided virtual storage if possible. If not enough virtual storage is provided for an internal sort, two external sort workfiles are built and the key-pointer pairs are sorted externally. If the minimum amount of virtual storage is not provided (see "How an Alternate Index is Built"), the BLDINDEX processing terminates with an error message.

When you specify EXTERNALSORT, or if not enough virtual storage is provided for an internal sort, you must provide two DD statements that describe the external-sort workfiles to be defined by BLDINDEX.

You can name the DD statements IDCUT1 and IDCUT2. When you choose other names for the workfile DD statements, you must identify the DD statements with the WORKFILES parameter.

Abbreviations: ESORT and ISORT

WORKFILES(dname b dname)

names the DD statements that identify the workfiles you want BLDINDEX to define if an external sort of the key-pointer pairs is required. You can use DD statements to describe two workfiles that will be defined and opened before the BLDINDEX routine begins processing the base-cluster's data records.

When you code the DD statements that describe the workfiles and identify them with the standard dnames IDCUT1 and IDCUT2, you do not need to specify the WORKFILES parameter.

For information on how to code the DD statements that describe the workfiles, see "DD Statements that Describe the SORT Workfiles."

Abbreviation: WFILE

BLDINDEX EXAMPLE

This example builds an alternate index over a previously-defined base cluster, EXAMPLE.KSDS2. Data records have already been loaded into the base cluster, so that it is not empty. The alternate index, its path, and its base cluster are all defined in the same catalog, AMASTCAT.

```
//BUILDAIX JOB ...
//STEP1 EXEC PGM=IDCAMS
//BASEDD DD DSN=EXAMPLE.KSDS2,DISP=OLD
//AIXDD DD DSN=EXAMPLE.AIX,DISP=OLD
//IDCUT1 DD DSN=SORT.WORK.ONE,DISP=OLD,AMP='AMORG',
// VOL=SER=VSER01,UNIT=2314
//IDCUT2 DD DSN=SORT.WORK.TWO,DISP=OLD,AMP='AMORG',
// VOL=SER=VSER01,UNIT=2314
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
```

```
BLDINDEX INFILE(BASEDD) -
        OUTFILE(AIXDD/AIXUPPW) -
        CATALOG(AMASTCAT/MCATUPPW)
```

/*

The job control language statements are:

- BASEDD DD, which describes the base cluster.
- AIXDD DD, which describes the alternate index.
- IDCUT1 and IDCUT2 DD, which describe a volume containing VSAM data space to be made available to BLDINDEX for defining and using two sort work data sets in the event an external sort is performed. This data space will not be used by BLDINDEX if enough virtual storage is available to perform an internal sort.

The BLDINDEX command builds an alternate index. The assumption is made that enough virtual storage will be available to perform an internal sort. However, note that DD statements with the default dnames of IDCUT1 and IDCUT2 have been provided for two external sort work data sets in the event that the assumption is incorrect and an external sort must be performed. The BLDINDEX command's parameters are:

- INFILE, which names the base cluster. The dname of the DD statement for this object must be identical to this name. Note that a password is not required since the base cluster is not protected.
- OUTFILE, which names the alternate index. The dname of the DD statement for this object must be identical to this name. The update password of the alternate index is also required.
- CATALOG, which identifies the master catalog. If it is necessary for BLDINDEX to use external sort work data sets, they will be defined in and deleted from the master catalog. The update password will permit these actions.

CHKLIST

The format of the CHKLIST command is:

CHKLIST	INFILE(dname) [CHECKID(checkid...)] [OUTFILE(dname)]
----------------	---

CHKLIST can be abbreviated: CKLST.

CHKLIST PARAMETERS**Required Parameter****INFILE(dname)**

specifies the name of the DD statement that identifies the checkpoint data set that contains the checkpoint entries to be processed.

Abbreviation: IFILE

Optional Parameters**CHECKID(checkid...)**

specifies one or more checkpoint identifiers of entries in the checkpoint data set for which to list tape data sets that were open at the time of the checkpoint. Each checkid must be one to sixteen characters in length. It must be enclosed in single quotation marks if it contains commas, semicolons, blanks, parentheses, slashes, or asterisks. The checkid must be the same as was specified in the CHKPT macro. A maximum of 255 checkids can be specified. Multiple checkids may be specified in any sequence.

If CHECKID is omitted, identification of tape data sets that were open at the time of a checkpoint are listed for all entries in the checkpoint data set. CHECKID must be omitted when a member of a partitioned checkpoint data set has been specified after dsname in the DD statement identified by INFILE.

If the checkpoint data set contains duplicates of a checkid, you can cause all of the checkpoint entries with that checkid to be listed either by specifying the checkid at least twice or by specifying, along with the checkid, a checkid for which there is no entry in the checkpoint data set. If you do neither, only the first checkpoint entry found with the checkid is listed.

Abbreviation: CHKID

OUTFILE(dname)

specifies the dname of the DD statement that identifies a data set other than the SYSPRINT data set to be used as an output data set. An output data set must meet the requirements shown under "Output Data Sets" in the "Introduction." If OUTFILE is not specified, the tape data set information is listed in the SYSPRINT data set.

Abbreviation: OFILE

CHKLIST EXAMPLES

Selecting Specific Checkpoint Entries: Example 1

In this example, the tape data sets that were open at checkpoint time are identified and listed on SYSPRINT for checkpoint entries C0000001 and C0000002.

```
//CHKLIST1 JOB      ...
//STEP1   EXEC     PGM=IDCAMS
//SYSPRINT DD      SYSOUT=A
//CHKPT   DD       DSN=CHKPT.DATASET,DISP=OLD
//SYSIN   DD       *
          CHKLIST -
          INFILE(CHKPT) -
          CHECKID(C0000001 C0000002)
/*
```

The job control statements are:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are sent.
- CHKPT DD, which identifies the checkpoint data set that contains the entries for which the tape data set information is to be listed.

The CHKLIST command prints the tape data set information as specified by the command's parameters:

- INFILE, which points to the CHKPT DD statement. The CHKPT DD statement identifies the checkpoint data set.
- CHECKID, which specifies that the checkpoint entries with checkids C0000001 and C0000002 are the only ones on the checkpoint data set for which the tape data set information is to be listed.

The output shown in "Appendix E. Sample Output From CHKLIST" was obtained with this JCL.

Partitioned Checkpoint Data Set: Example 2

In this example, the checkpoint data set is a partitioned data set, and member C0000001 is selected for processing by CHKLIST.

```
//CHKLIST2 JOB      ...
//STEP1   EXEC     PGM=IDCAMS
//SYSPRINT DD      SYSOUT=A
//CHKPTDD DD       DSN=EXAMPLE.CHKPTDS2(C0000001),
//         DISP=SHR
//SYSIN   DD       *
          CHKLIST -
          INFILE(CHKPTDD)
/*
```

The job control language statements are:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.
- CHKPTDD DD, which identifies the partitioned checkpoint data set and the specific member for which the tape data set information is to be listed.

The CHKLIST command prints the tape data set information as specified by its parameter, INFILE, which points to the CHKPTDD DD statement. The CHKPTDD DD statement identifies the member of the partitioned checkpoint data set.

Note: The CHECKID parameter is omitted so that the checkpoint entry C0000001 will automatically be processed by CHKLIST.

CNVTCAT

The format of the CNVTCAT command is:

CNVTCAT	{INFILE(dname) INDATASET(entryname)} [CATALOG(catname[/password])] [CVOLEQUATES((catname b(volser [b volser...]))[b(catname...)...])] [LIST NOLIST] [MASTERCATALOG(catname[/password])]
---------	--

CNVTCAT can be abbreviated: CNVTC

CNVTCAT PARAMETERS

Required Parameters

INFILE(dname)
specifies the name of a DD statement that identifies the OS catalog that is to be converted.

Abbreviation: IFILE

INDATASET(entryname)
specifies the control-volume-pointer (CVPE) name (not the data set name) of the OS catalog that is to be converted. If INDATASET is specified, the OS catalog is dynamically allocated. The OS system catalog entryname, SYSCTLG, cannot be specified. If entries in the system catalog are being converted, INFILE must be used to specify a DD statement that identifies the OS system catalog.

Abbreviation: IDS

Optional Parameters

CATALOG(catname[/password])
specifies the name of a catalog that is to receive the converted entries. If CATALOG is not specified, see "Order of Catalog Use: CNVTCAT" for information about the order in which a catalog is selected.

password
specifies, for a password-protected catalog, the update- or higher-level password.

Abbreviation: CAT

CVOLEQUATES((catname b(volser[b volser...]))[b(catname...)])
identifies the user catalog that converted control-volume-pointer entries (CVPEs) point to. CVPEs point to control volumes (CVOLs) which contain OS catalog entries. The OS catalog entries in a CVOL might have been (or might soon be) converted to VSAM catalog entries and put into a user catalog. When a CVPE is converted to an alias entry, the alias entry points to the user catalog that contains (or is to contain) the control volume's converted OS catalog entries.

When you specify CVOLEQUATES, you must also specify MASTERCATALOG.

catname
specifies the name of an existing VSAM catalog. The catalog named contains or is to contain converted OS catalog entries that were cataloged in the CVOL pointed to by the CVPE being converted.

volser

specifies the volume serial number(s) of one or more CVOLS for which entries have been or are to be converted.

Abbreviation: CVEQU

LIST|NOLIST

specifies whether entries are to be listed after they are converted.

Abbreviation: NLST

MASTERCATALOG(catname[/password])

specifies the name of the master catalog into which any aliases for user catalogs are to be placed. MASTERCATALOG is required when CVOLEQUATES is specified.

password

specifies the master catalog's update- or higher-level password.

Abbreviation: MCAAT

CNVTCAT EXAMPLES

The two CNVTCAT examples are related and show how the entries of two catalogs in an OS system can be converted to entries in a VSAM catalog. The catalog on volume VSER09 contains control-volume-pointer entries (CVPEs) that point to the catalog on volume VSER08. Therefore, the two catalogs are chained together with VSER08's catalog at the end of the chain. When two or more OS catalogs are chained together, the catalog at the end of the chain should be converted first.

Convert an OS Catalog's Entries to Entries in a VSAM Catalog: Example 1

In this example, the entries in the catalog on volume VSER08 are converted to VSAM catalog entries and written into the USER11 catalog.

```
//CNVTCAT1 JOB      ...
//STEP1    EXEC    PGM=IDCAMS
//SYSPRINT DD      SYSOUT=A
//OSCAT1   DD      VOL=SER=VSER08,DISP=OLD,DSN=SYSCTLG,
//          UNIT=2314
//SYSIN    DD      *
              CNVTCAT -
                INFILE(OSCAT1) -
                CATALOG(USER11/MR)
/*
```

The job control language statements are:

- OSCAT1 DD, which describes and allocates the OS catalog that is to be converted.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The CNVTCAT command converts the entries in the OS catalog on volume VSER08 to entries in a user catalog, USER11. The command's parameters are:

- INFILE, which points to the OSCAT1 DD statement. The OSCAT1 DD statement describes and allocates the OS catalog that is to be converted.
- CATALOG, which specifies the name of an existing VSAM catalog, the user catalog USER11. USER11 is to receive the converted entries. The user catalog's update- (or higher-level) password, MR, is required.

Convert an OS Catalog's Entries to Entries in a VSAM Catalog: Example 2

In this example, the entries in the OS catalog on volume VSER09 are converted to VSAM catalog entries and written into the USER12 catalog. Some of the entries in the OS catalog are control volume pointer entries, and point to the OS catalog on volume VSER08 (which was converted in the previous example).

```
//CNVTCAT2 JOB      ...
//STEP1   EXEC     PGM=IDCAMS
//SYSPRINT DD      SYSOUT=A
//OSCAT2  DD        VOL=SER=VSER09,DISP=OLD,DSN=SYSCTLG,
//        UNIT=2314
//SYSIN   DD        *
          CNVTCAT -
          INFILE(OSCAT2) -
          CATALOG(USER12/MR) -
          CVOLEQUATES( -
            (USER11 (VSER08)) ) -
          NOLIST -
          MASTERCATALOG(AMASTCAT)
/*
```

The job control language statements are:

- OSCAT2 DD, which describes and allocates the OS catalog that is to be converted.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The CNVTCAT command converts the entries in the OS catalog on volume VSER09 to entries in a user catalog, USER12. Because CVOLEQUATES is specified, an alias entry is built and put in the master catalog to relate the name VSER08 to the user catalog USER11. The command's parameters are:

- INFILE, which points to the OSCAT2 DD statement. The OSCAT2 DD statement describes and allocates the OS catalog to be converted.
- CATALOG, which specifies the name of an existing VSAM catalog, the user catalog USER12. USER12 is to receive the converted entries. The user catalog's update- (or higher-level) password, MR, is required.
- CVOLEQUATES, which specifies that control-volume-pointer entries (CVPEs) are to be converted to alias entries. Each CVPE in VSER09's catalog points to an entry in VSER08's catalog. Each new alias entry points to the USER11 user catalog connector entry in the master catalog. The alias entries are written into the master catalog.
- NOLIST, which specifies that the entries are not to be listed after they are converted.
- MASTERCATALOG, which identifies the master catalog. This parameter must be specified because CVOLEQUATES is also specified. If the master catalog is password protected, the catalog's update- (or higher-level) password must be specified.

DEFINE ALIAS

The format of the DEFINE command when it is used to define an alias is:

DEFINE	ALIAS (NAME(aliasname) RELATE(entryname)) [CATALOG(catname[/password])]
---------------	---

DEFINE may be abbreviated: DEF.

DEFINE ALIAS PARAMETERS

Required Parameters

ALIAS

specifies that an alias for a user catalog or nonVSAM data set is to be defined.

NAME(aliasname)

specifies the alias—the alternate entryname for a user catalog or nonVSAM data set.

aliasname can contain 1 to 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and special characters (the hyphen (-), and the 12-0 overpunch (X'C0')).

When the name contains more than eight characters, you divide the name into segments of 1 to 8 characters each and separate the segments with (.) periods.

The name's first character (and each segment's first character) is either an alphabetic character or a national character.

RELATE(entryname)

specifies the name of the entry (that is, the user-catalog entryname or the nonVSAM data set name) for which the alias is being defined.

An alias entry must reside in the same catalog as the entry to which it is related. When an alias is defined, catalogs are searched for the related entry. When the related entry is found, the alias entry is added to the catalog. An alias must be unique within a catalog. An alias can be defined only for a nonVSAM data set or for a user-catalog connector in the master catalog.

Abbreviation: REL

Optional Parameter

CATALOG(catname[/password])

identifies the catalog in which the alias is to be defined. If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved. See "Order of Catalog Search: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

catname

specifies the name of the catalog. When the alias is for a user catalog connector, catname is the name of the master catalog.

password
specifies a password. If the catalog is password protected, you must supply the catalog's update- or higher-level password. If no password is specified and the catalog is password protected, VSAM asks the operator or TSO terminal user for the correct password.

Abbreviation: CAT

DEFINE ALIAS EXAMPLES

Define an Alias for a NonVSAM Data Set: Example 1

In this example, an alias is defined for a nonVSAM data set.

```
//DEFALS JOB ...
//STEP1 EXEC PGM=IDCAMS
//STEP1 DD DSNAME=D27UCAT1,DISP=OLD
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DEFINE ALIAS -
(NAME(EXAMPLE.NONVSAM1) -
RELATE(EXAMPLE.NONVSAM) )
/*
```

The job control language statements are:

- STEPCAT DD, which makes a catalog available for this job step: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE ALIAS command defines an alias, EXAMPLE.NONVSAM1, for the nonVSAM data set EXAMPLE.NONVSAM. Its parameters are:

- NAME, which specifies the alias (alternate entryname), EXAMPLE.NONVSAM1.
- RELATE, which specifies the name that the alias is an alternate entryname for, EXAMPLE.NONVSAM.

Define an Alias for a User Catalog: Example 2

In this example, an alias is defined for a user catalog. The alias is defined in the master catalog.

```
//DEFUCALS JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DEFINE ALIAS -
(NAME(D40) -
RELATE(D27UCAT1)) -
CATALOG(AMASTCAT/MCATUPPW)
/*
```

The job control statements are:

- SYSPRINT, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE ALIAS command defines an alias, D40, for the VSAM user catalog, D27UCAT1. Any data set defined in D27UCAT1 having 'D40' as its first qualifier will be located by VSAM in D27UCAT1 when it is referenced in an Access Method Services command or user program without requiring a JOBCAT or STEPCAT DD statement for D27UCAT1. D27UCAT1 can also be referenced using the alias D40. The parameters are:

- NAME, which specifies the alias, D40.
- RELATE, which specifies the name of the user catalog for which D40 is an alternate entryname (alias).
- CATALOG, which supplies the update password of the master catalog.

DEFINE ALTERNATEINDEX

When you define an alternate index, you can specify attributes for the alternate index as a whole and for the alternate-index's components. The format of the DEFINE ALTERNATEINDEX command is:

DEFINE ALTERNATEINDEX (parameters) -

[DATA(parameters)] -

[INDEX(parameters)] -

[CATALOG(subparameters)]

DEFINE ALTERNATEINDEX

DEFINE	<pre> ALTERNATEINDEX (NAME(entryname) RELATE(entryname[/password]) {CYLINDERS(primary[b secondary]) RECORDS(primary[b secondary]) TRACKS(primary[b secondary])} VOLUMES(volser[b volser...]) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [DESTAGWAIT <u>NODESTAGWAIT</u>] [ERASE <u>NOERASE</u>] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE(CI-percent[b CA-percent] 0 0)] [IMBED <u>NOIMBED</u>] [KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])] [KEYS(length b offset 64 0)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [ORDERED <u>UNORDERED</u>] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE(average b maximum 4086 32600)] [REPLICATE <u>NOREPLICATE</u>] [REUSE <u>NOREUSE</u>] [SHAREOPTIONS(crossregion [b crosssystem] 1 3)] [SPEED <u>RECOVERY</u>] [STAGE <u>BIND</u> CYLINDERFAULT] [TO(date) FOR(days)] [UNIQUE <u>SUBALLOCATION</u>] [UNIQUEKEY <u>NONUNIQUEKEY</u>] [UPDATEPW(password)] [UPGRADE <u>NOUPGRADE</u>] [WRITECHECK <u>NOWRITECHECK</u>] </pre>
--------	--

(Continued on next page.)

	<pre> [DATA ([ATTEMPTS(number)) [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [CYLINDERS(primary[b secondary])] RECORDS(primary[b secondary])] TRACKS(primary[b secondary])] [DESTAGEWAIT NODESTAGEWAIT] [ERASE NOERASE] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE(CI-percent[b CA-percent])] [KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])] [KEYS(length b offset)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [NAME(entryname)] [ORDERED UNORDERED] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE(average b maximum)] [REUSE NOREUSE] [SHAREOPTIONS(crossregion[b crosssystem])] [SPEED RECOVERY] [STAGE BIND CYLINDERFAULT] [UNIQUE SUBALLOCATION] [UNIQUEKEY NONUNIQUEKEY] [UPDATEPW(password)] [VOLUMES(volser[b volser...])] [WRITECHECK NOWRITECHECK]]] [INDEX ([ATTEMPTS(number)) [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [CYLINDERS(primary[b secondary])] RECORDS(primary[b secondary])] TRACKS(primary[b secondary])] [DESTAGEWAIT NODESTAGEWAIT] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [IMBED NOIMBED] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [NAME(entryname)] [ORDERED UNORDERED] [OWNER(ownerid)] [READPW(password)] [REPLICATE NOREPLICATE] [REUSE NOREUSE] [SHAREOPTIONS(crossregion[b crosssystem])] [STAGE BIND CYLINDERFAULT] [UNIQUE SUBALLOCATION] [UPDATEPW(password)] [VOLUMES(volser[b volser...])] [WRITECHECK NOWRITECHECK]]] [CATALOG(catname[/password])] </pre>
--	---

DEFINE ALTERNATEINDEX Parameters: Summary

The parameters of the DEFINE ALTERNATEINDEX command can be grouped as required and optional parameters.

The required parameters are:

- Entry type, which specifies that an alternate index and its catalog entry are to be created (ALTERNATEINDEX).
- Name, which names the alternate index (NAME).
- Base cluster, which identifies the base cluster related to the alternate index (RELATE).
- Allocation, which specifies the amount of space to be allocated to the alternate index and identifies the volume(s) that is to contain the alternate index (VOLUMES and one of CYLINDERS, RECORDS, or TRACKS). You need not explicitly specify the allocation parameters if you use the MODEL parameter.

The optional parameters are:

- Model, which identifies an existing alternate index entry that is to be used as a model for the alternate index being defined (MODEL).
- Allocation, which specifies
 - the name of the DD statement that describes the alternate-index's volume(s) (FILE).
 - the length and offset of the alternate-key field in the base cluster (KEYS).
 - whether the alternate-index records are to be divided by alternate-key value among volumes (KEYRANGES).
 - whether the alternate-key value can identify more than one data record in the base cluster (UNIQUEKEY or NONUNIQUEKEY).
 - the size of control intervals, records, and the user program's buffer area (CONTROLINTERVALSIZE, RECORDSIZE, and BUFFERSPACE).
 - how an alternate index or component that is stored on a mass storage device is to be staged (STAGE, BIND, or CYLINDERFAULT).
 - the amount of free space to be left in control intervals and control areas when the alternate-index records are loaded (FREESPACE).
 - whether the alternate index is to share its data space with other VSAM objects (SUBALLOCATION or UNIQUE).
 - whether the volumes are to be allocated in the order given with the VOLUMES parameter (ORDERED or UNORDERED).
 - whether the alternate index is to be temporary and reusable (REUSE or NOREUSE).
- Data integrity, which specifies
 - the name of a user-written I/O error routine (EXCEPTIONEXIT).
 - whether alternate-index records are to be overwritten with zeros when the alternate index is deleted (ERASE or NOERASE).

- whether the alternate index can be shared among subtasks, jobs, and systems (SHAREOPTIONS).
- whether the alternate-index records are to be checked when the direct access device writes them into secondary storage (WRITECHECK or NOWRITECHECK).
- whether an alternate index or component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that closes it (DESTAGWAIT or NODESTAGWAIT).
- whether space is to be preformatted when the alternate index is initially loaded (RECOVERY or SPEED).
- whether VSAM is to upgrade the alternate index each time its base cluster is modified (UPGRADE or NOUPGRADE).
- Protection, which
 - specifies passwords for the alternate index (MASTERPW, CONTROLPW, UPDATEPW, and READPW).
 - identifies the alternate index to a console operator without revealing its entryname (ATTEMPTS and CODE).
 - identifies a user-written-security-verification routine for additional authorization verification (AUTHORIZATION).
 - identifies the alternate index's owner (OWNER).
 - specifies a retention period for the alternate index (FOR or TO).
- Parameters that apply to the index of an alternate index, which specify
 - whether index records are to be repeated around the track of the index's direct access device (REPLICATE or NOREPLICATE).
 - whether the index's sequence-set records are to be written adjacent to their data control areas (IMBED or NOIMBED).
- Components, which specifies attributes for the alternate-index's data and index components that might differ from the attributes specified for the alternate index as a whole or the attributes established for the alternate index being used as a model (DATA and INDEX).
- Catalog, which identifies the catalog in which the alternate index is to be defined (CATALOG).

DEFINE can be abbreviated DEF.

DEFINE ALTERNATEINDEX PARAMETERS

Required Parameters

ALTERNATEINDEX

specifies that an alternate index is to be defined and have space allocated.

The ALTERNATEINDEX keyword is followed by the parameters specified for the alternate index as a whole. These parameters are enclosed in parentheses and, optionally, are followed by the DATA and INDEX parameters and their subparameters.

Abbreviation: AIX

NAME(entryname)

specifies the alternate-index's entryname or the name of each of its components. The entryname specified for the alternate index as a whole is not propagated to the alternate-index's components. You can define a separate entryname for the alternate index, its data component, and its index component. If no name is specified for the data and index component, a name is generated. Because the alternate index, data component, and index component are individually named, each can be addressed.

NAME is a required parameter of ALTERNATEINDEX. When NAME is not specified as a parameter of DATA or INDEX, VSAM generates an entryname for the alternate-index's components.

entryname

can contain 1 to 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and special characters (the hyphen (-), and the 12-0 overpunch (X'C0')).

When the name contains more than eight characters, you divide the name into segments of 1 to 8 characters each and separate the segments with (.) periods. The name's first character (and each segment's first character) is either an alphabetical character or a national character. If you specify the name in the format entryname(modifier), only that portion of the name preceding the left parenthesis will be used. The modifier enclosed in parentheses will be ignored.

RELATE(entryname[/password])

names the alternate-index's base cluster. The base cluster is an entry-sequenced cluster or a key-sequenced cluster that the alternate index is to be organized over.

You can relate an alternate index only to a key-sequenced cluster or to an entry-sequenced cluster. You cannot relate an alternate index to a reusable cluster or to a relative-record cluster.

entryname

names the base cluster.

password

specifies the base cluster's master password.

You must specify the base cluster's master password when the base cluster is password protected. If you do not specify the base cluster's master password, you can specify the catalog's master password (that is, the master password of the catalog that contains the base cluster's entry—see the CATALOG parameter). If no password is specified with either the RELATE or the CATALOG parameter, VSAM asks the operator for the base cluster's master password.

Abbreviation: REL

CYLINDERS(primary[b secondary])

RECORDS(primary[b secondary])

TRACKS(primary[b secondary])

specifies the amount of space to be allocated to the alternate index from the volume's available space, when UNIQUE is specified, or from the available space of one of the volume's VSAM data spaces. If you specify RECORDS, the amount of space allocated is the minimum number of tracks that are required to contain the specified number of records. However, if you specify TRACKS or RECORDS and if the minimum number of tracks should exceed a cylinder, space is allocated in terms of cylinders.

When UNIQUE is specified and the data component is to be contained on more than one volume, each volume can contain a maximum of 16 extents.

When the data component is not divided into key ranges and more than one volume is specified, the primary amount of space is allocated only on the first volume when the component is defined. When the component increases in size so as to extend on to additional volumes, the first allocation on each overflow volume is the primary amount.

When KEYRANGES is specified and the data component is to be contained on more than one volume, the primary amount of space is immediately allocated on each volume required for the key ranges. If more than one key range is allocated on a single volume, the primary amount is allocated for each key range.

secondary amounts can be allocated on all volumes available to contain parts of the alternate index regardless of the key ranges when the alternate index is extended.

The primary and each secondary allocation must be able to be satisfied within 5 extents; otherwise, your DEFINE or data set extension will fail.

If you do not specify the MODEL parameter, you must specify one, and only one, of the following parameters: CYLINDERS, RECORDS, or TRACKS. You can specify the amount of space as a parameter of ALTERNATEINDEX, as a parameter of DATA, or as a parameter of both DATA and INDEX. If the space is specified as a parameter of:

- ALTERNATEINDEX, the amount specified is divided between the data and index components. The division algorithm is a function of control interval size, record size, device type, and other data set attributes. If the division results in an allocation for the data component which is not an integral multiple of the required control-area size, the data component's allocation is rounded up to the next higher control-area multiple. This rounding may result in a larger total allocation for your alternate index than that which you specified.
- DATA, the entire amount specified is allocated to the data component. An additional amount of space, depending on control interval size, record size, device type, and other data set attributes, is allocated to the index component.

When you specify UNIQUE, and the amount of space you specify for the data component of the alternate index is not an even multiple of cylinder size, VSAM rounds the allocation upward to a multiple of a cylinder.

To determine the exact amount of space allocated to each component, list the alternate-index's catalog entry, using the LISTCAT command.

When you specify UNIQUE and the alternate-index's data space is the first data space on a volume that belongs to a recoverable catalog, one additional cylinder is allocated for the catalog recovery area.

primary

specifies the initial amount of space that is to be allocated to the alternate index.

secondary

specifies the amount of space that is to be allocated each time the alternate index extends, as a secondary extent. When secondary is specified, space for the alternate-index's data and index components can be expanded to include a maximum of 123 extents.

If the data space that contains (all or part of) the cluster cannot be extended because the cluster's secondary allocation amount is greater than the data-space's secondary allocation amount, VSAM builds a new data space. The new data space's primary and secondary allocation amounts are equal to the cluster's secondary allocation amount.

primary and secondary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviations: CYL, REC, and TRK

VOLUMES(volser[*b* volser...])

specifies the volume(s) on which an alternate-index's components are to have space.

If you do not specify the MODEL parameter, VOLUMES must be specified either as a parameter of ALTERNATEINDEX, or as a subparameter of both DATA and INDEX. If the data and index components are to reside on different device types, you must specify VOLUMES as a parameter of both DATA and INDEX.

If more than one volume is listed with a single VOLUMES parameter, the volumes must be of the same device type. The volume serial number may be repeated in the list only if the KEYRANGE parameter is specified. You may wish to do this in order to have more than one keyrange on the same volume. Even in this case, repetition is only valid if all duplicate occurrences are used for the primary allocation of some keyrange.

In a system with the Mass Storage System, an alternate index or component can be defined on a mass storage volume.

volser

is a 1 to 6 alphameric or special character volume serial number.

If the volser contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the volser contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the volser in single-quotation marks (for example, VOLUMES('*DORIS')).

When the volser contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, VOLUMES('*CA' 'ROL')).

You can code volser in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, VOLUMES(X'E2E4C5') is the same as VOLUMES(SUE).

The VOLUMES parameter interacts with other DEFINE ALTERNATEINDEX parameters. You should take care to ensure that the volume(s) you specify for the alternate index can satisfy the alternate index's other attributes:

- **SUBALLOCATION:** If **UNIQUE** is not specified, the volume must already contain VSAM data space.
- **KEYRANGES:** If **KEYRANGES** and **UNIQUE** are specified, a VSAM data space is built and allocated on a separate volume for each key range.
- **ORDERED:** If **ORDERED** is specified, the volumes are used in the order listed. If **ORDERED** and **KEYRANGES** are specified, there is a one-for-one correspondence between the key ranges in the key range list and the volumes in the volser list.
- **CYLINDERS, RECORDS, TRACKS:** The volume(s) contains enough available space to satisfy the component's primary space requirement.
- **FILE:** When **UNIQUE** is specified for the alternate index, the volume information supplied with the DD statement pointed to by **FILE** is consistent with the information specified for the alternate index and its components.
- **CATALOG:** If the volume already has a data space on it, subsequent data spaces can be defined for the alternate index and must be cataloged in the catalog that owns the volume.

Abbreviation: VOL

Optional Parameters

ATTEMPTS(number|2)

specifies the maximum number of times the operator can try to enter a correct password in response to a prompting message.

This parameter can be coded, but only has effect when the alternate-index's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

number

is an integer from 0 to 7 and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

When **ATTEMPTS(0)** is coded, the operator is not prompted and is not allowed to enter a password from the console.

Note to TSO users: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password for the alternate index. Checking the logon password counts as one attempt to obtain a password. If **ATTEMPTS** is not specified, the user has one attempt to supply the alternate-index's password because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])

specifies that a user-security-verification routine (USVR) is available for additional security verification. When a protected alternate index is opened and the user supplies a correct password other than the alternate-index's master password, the USVR receives control. For details on coding a USVR, see the OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.

If a USVR is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

This parameter can always be coded, but only has effect when the master password is not null.

entrypoint

specifies the name of the USVR. entrypoint can contain 1 to 8 alphanumeric characters (A through Z and 0 through 9) or national characters (@, #, \$), and special characters (the hyphen (-) and the 12-0 overpunch (X'CO')). The name's first character is either an alphabetic character or a national character.

string

specifies information to be passed to the USVR when it receives control to verify authorization. string can contain 1 to 255 EBCDIC characters.

When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single-quotation marks (for example, AUTHORIZATION(entrypoint, 'DORIS')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, 'CA'ROL')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint, SUE). The string can contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

BUFFERSPACE(size)

specifies the minimum space that your program's address space is to provide for buffers. The buffer space size you specify helps VSAM determine the data component's and index component's control interval size.

If the specified size is less than VSAM requires for the buffers needed to run your job, VSAM terminates your job and provides an appropriate error message.

When you do not specify BUFFERSPACE, VSAM determines the buffer space size. VSAM determines the control interval size first, then sets the buffer space amount equal to two data control intervals and one index control interval.

When you specify BUFFERSPACE, you must specify at least enough space to contain two data control intervals and one index control interval.

size

is the amount of space to be provided for buffers. The value can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviations: BUFSPC or BUFSP

CODE(code)

specifies a code name for the alternate index. If an attempt is made to access a password-protected alternate index without first supplying an appropriate password, a prompting message might be issued to the operator's console (see ATTEMPTS above). The prompting message includes the code name, which identifies the alternate index without revealing its entryname.

This parameter can be coded, but only has effect when the alternate-index's master password is not null. A prompting message is issued only when the user hasn't already supplied the appropriate password.

When you don't specify a code name for the alternate index, the prompting message identifies the alternate index with its entryname.

code

can contain 1 to 8 alphameric or special characters. When the code contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the code in single-quotation marks (for example, CODE('*DORIS*')).

If the code contains a lowercase alphabetic character it will be folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE('*CA'ROL*')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CODE(X'E2E4C5') is the same as CODE(SUE).

CONTROLINTERVALSIZE(size)

specifies the size of the alternate-index's control intervals. The size of the control interval depends on the maximum size of data records, and on the amount of buffer space specified.

When you do not specify the control interval size, VSAM determines the control interval size. If you have not specified BUFFERSPACE and the size of your records permits, VSAM selects the optimum size for the data control interval size and 512 bytes for the index control interval size.

size

for the alternate-index's data component.

Since an alternate index always has the spanned attribute, the control interval size can be less than the maximum record length. The sizes you can specify—between 512 and 32,768 bytes—are a multiple of 512 or 2048:

$$CISZ = (n \times 512) \text{ or } (n \times 2048)$$

where n is a positive integer from 1 to 16

size

for the alternate-index's index component.

You can specify the following values for the index component's control interval size:

$$CISZ = [512 \mid 1024 \mid 2048 \mid 4096]$$

When you specify a size that is not a multiple of 512 or 2048, VSAM chooses the next higher multiple.

Refer to the chapter "Optimizing VSAM's Performance" in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for a discussion of the relationship between control interval size and physical block size. The discussion also includes restrictions that apply to control interval size and physical block size.

Abbreviations: CISZ or CNVSZ

CONTROLPW(password)

specifies a control password for the alternate index and its data and index components. The control password allows read and write operations against the object's control intervals.

When specified as a parameter of DATA or INDEX, the control password allows the user's program to open the data component or index component for read and write processing of the component's control intervals (that is, the entire control interval, including the data portion of stored records and the control fields VSAM inserts into stored records and control intervals).

If a read or update password is the only password specified for the object, it (the highest-level password) propagates upward and becomes the password for all higher unspecified levels.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

DESTAGEWAIT | NODESTAGEWAIT

specifies whether an alternate index or its component that is stored on a mass storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it.

DESTAGEWAIT indicates that destaging is to be completed before VSAM returns control to the program that issued the CLOSE macro. VSAM can thus notify the program whether destaging was successful.

NODESTAGEWAIT indicates that notification of unsuccessful destaging is to be made only by a message to the operator.

When the alternate index or component isn't stored on a mass storage volume, the attribute is ineffective until the direct access storage volume the alternate index or component is stored on is converted to a mass storage volume (by way of the CONVERTV command, which is described in OS/VS Mass Storage System (MSS) Services for Reference Information).

When one of these parameters is specified for the data component and the other parameter is specified for the index component, the components are destaged separately as

specified, except when the sequence set of the index component is imbedded in the data. In that case, the parameter specified for the index component applies to both components.

Abbreviations: DSTGW and NDSTGW

ERASE|NOERASE

specifies whether the alternate index data component's records are to be erased when the alternate index is deleted. If ERASE is specified, the alternate index data component's records are overwritten with binary zeros when the alternate index is deleted. This attribute applies only to the alternate-index's data component.

Abbreviations: ERAS and NERAS

EXCEPTIONEXIT(entrypoint)

specifies the name of the user-written routine, called the exception exit routine, that receives control when an exceptional I/O error condition occurs during the transfer of data between your program's address space and the alternate-index's direct access storage space. (An exception is any condition that causes a SYNAD exit to be taken.) The component's exception-exit routine is processed first, then the user's SYNAD exit routine receives control.

If an exception exit routine is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

Abbreviation: EEXT

FILE(dname)

names the DD statement that identifies the direct access device(s) and volume(s) on which space is to be allocated to the alternate index. When you specify UNIQUE, FILE can identify the DD statement which identifies and allocates the volumes of the alternate index. If more than one volume is specified in a volume list, all volumes must be the same device type. When the data component and index component are to reside on separate devices, you can specify a separate FILE parameter as a subparameter of DATA and INDEX to point to different DD statements.

When the alternate index is defined in a recoverable catalog, and FILE is specified as a parameter of ALTERNATEINDEX, the FILE parameter can identify the DD statement for all volumes on which space is to be allocated. You also use FILE to identify the volume containing the base cluster's index component (when the base cluster is key-sequenced) or data component (when the base cluster is entry-sequenced) when defining into a recoverable catalog. When FILE refers to more than one volume of the same device type, the DD statement that describes the volumes cannot be a concatenated DD statement. When the alternate index is cataloged in a recoverable catalog, the DD statement can be concatenated if the volumes are a different device type: part of the concatenated DD statement describes the alternate index's or component's volumes (of one device type); the other part of the DD statement describes the volume that contains the catalog recovery area (of another device type).

If the FILE parameter is not specified, the required volumes are dynamically allocated. If the alternate index or one of its components is unique, and if the FILE parameter is not specified and the volume(s) is physically mounted, the volume(s) identified with the VOLUMES parameter is dynamically allocated. The volume(s) must be mounted as permanently resident or reserved.

The DD statement you specify must be in the form:

```
//ddname DD UNIT=(devtype[,unitcount]),  
// VOL=SER=(volser1, volser2, volser3, ...), ...
```

FREESPACE(CI-percent[b CA-percent]|0 0)

specifies the amount of space that is to be left empty after any allocation (that is, the initial, or primary, allocation and each extension, or secondary allocation) and any split of control intervals (CI-percent) and control areas (CA-percent) when the alternate index is built (see the BLDINDEX command). The amount of empty space in the control interval and control area is available for data records that are updated and inserted after the alternate index is initially built.

The amounts are specified as percentages. CI-percent translates into a number of bytes that is equal to, or slightly less than, the percentage value of CI-percent. CA-percent translates into a number of control intervals that is equal to, or less than, the percentage value of CA-percent. The percentages, which must be equal to or less than 100, can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

When you specify 100% of free space, one data record is placed in the first control interval of each control area when the alternate index is built.

Abbreviation: FSPC

IMBED|NOIMBED

specifies whether the sequence set (that is, the lowest level of the index of the alternate index—the level that points to alternate-index records) is to be put adjacent to the alternate-index's data component.

When you specify IMBED, the sequence-set record for each control area is written as many times as it will fit on the first track adjacent to the control area.

The IMBED|NOIMBED parameter interacts with the REPLICATE|NOREPLICATE parameter in determining the physical attributes of the index of the alternate index. These index attributes can be used in the following combinations:

- The sequence-set records are adjacent to the data control areas, and only the sequence-set records are replicated (IMBED and NOREPLICATE).
- The sequence-set records are adjacent to the data control intervals, and all levels of index records are replicated (IMBED and REPLICATE).
- All index records are together, and all index records are replicated (REPLICATE and NOIMBED).
- All index records are together, and no index records are replicated (NOREPLICATE and NOIMBED).

For some applications, specifying index options can improve the application's performance. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for information on how the index's optional attributes affect performance.

Abbreviations: IMBD and NIMBD

**KEYRANGES((lowkey b highkey)
[(lowkey b highkey)...])**

specifies that portions of the alternate-index's data component are to be put on different volumes. Each portion of the alternate index is called a key range.

The maximum number of alternate-key ranges is 123. Key ranges must be in ascending order, and aren't allowed to overlap. However, a gap can exist between two key ranges—you are not allowed to insert records within the gap.

Keys can contain 1 to 64 characters; 1 to 128 hexadecimal characters if coded as X'lowkey' b X'highkey'. All EBCDIC characters are allowed. Keys consisting of characters must be enclosed in single quotation marks if they contain commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in single quotation marks.

lowkey

specifies the low key of the key range. If lowkey is shorter than the actual keys, it will be padded on the right with binary zeros.

highkey

specifies the high key of the key range. If highkey is shorter than the actual keys, it will be padded on the right with binary ones.

The KEYRANGES parameter interacts with other DEFINE ALTERNATEINDEX parameters. You should take care to ensure that, when you specify KEYRANGES, the alternate-index's other attributes can be satisfied.

- **VOLUMES:** There should be as many volume-serial-numbers in the volser list as there are key ranges. When a volume serial number is duplicated in the volser list, more than one key range is allocated space on the volume. When more than one key range is contained on a volume, UNIQUE cannot be coded for the alternate-index's data component.

When there are more volumes in the volser list than there are key ranges, the excess volumes are used for overflow records from any key range without consideration for key range boundaries.

When there are fewer volumes in the volser list than there are key ranges, the excess key ranges are allocated on the last volume specified—UNIQUE cannot also be specified.

- **UNIQUE:** When UNIQUE is specified, each key range resides in its own VSAM data space on a separate volume. Other key ranges for the alternate index cannot also reside on the volume.
- **ORDERED:** There is a one-for-one correspondence between the volumes in the volser list and the key ranges: the first volume on the volume list contains the first key range, the second volume contains the second key range, and so on. If a volume can't be allocated in the order specified by the volser list, your alternate index definition job terminates with an error message.
- **KEYS:** the length of the key values must not exceed the key length specified in the KEYS parameter. The KEYS parameter must be specified in the same component as the KEYRANGE parameter.

Abbreviation: KRNG

KEYS(length b offset|64 0)

describes the alternate-key field in the base cluster's data record. The key field of a key-sequenced cluster is called the prime key, to distinguish it from other key fields, called alternate keys. The data record's alternate key can overlap or be contained entirely within another

(alternate or prime) key field.

The sum of length plus offset cannot be greater than the length of the base cluster's data record.

When the base cluster's data record is allowed to span control intervals, the record's alternate-key field is within the record's first segment (that is, in the first control interval.)

length b offset

specifies the length of the alternate key (between 1 and 255), in bytes, and its displacement from the beginning of the base cluster's data record, in bytes. You can express length and offset in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

MASTERPW(password)

specifies a master password for the alternate index. The master password allows all Access Method Services operations against the alternate index entry and its data and index entries. The master password allows the user's program to access the alternate-index's contents without restriction. For more details on how passwords can be used, see "Data Security and Protection."

If all passwords are null, ATTEMPTS, AUTHORIZATION, and CODE can be specified but have no effect until the master password is specified.

When specified as a parameter of DATA or INDEX, the master password allows the user's program to open the data component or index component and process it without restriction.

When you don't specify MASTERPW, but do specify another password, the object's highest-level password propagates upward and becomes the password for all higher levels, including the master password.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*').)

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA' 'ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation: MRPW

MODEL(entryname[/password]
[b catname[/password]])

identifies an existing alternate index entry that is to be used as a model for the alternate index being defined.

When you use an alternate index entry as a model for the alternate index, the model alternate-index's data and index entries are used as models for the to-be-defined alternate-index's data and index components, unless another entry is specified with the MODEL parameter as a subparameter of DATA or INDEX.

You can use an existing data entry as a model for the alternate-index's data component, and/or an existing index entry as a model for the alternate-index's index component, without specifying a model for the alternate index itself.

You can use some attributes of the model entry and override other attributes by explicitly specifying them in the alternate-index's definition. For details about how a model is used, see "How to Use One Object as a Model for Another Object." When you do not want to add or change any attributes, you specify only the entry type (ALTERNATEINDEX), the alternate-index's name (NAME), its base cluster's name (RELATE) and the model-entry's name (MODEL).

entryname

names the entry to be used as a model. entryname names an alternate index entry when MODEL is a parameter of ALTERNATEINDEX.

entryname names an alternate-index's data entry when MODEL is a parameter of DATA, and it names an alternate-index's index entry when MODEL is a parameter of INDEX.

password

specifies a password. If the model entry is password protected and it is cataloged in a password-protected catalog, you supply the read- (or higher-level) password of either the model entry or its catalog. If you supply both passwords, the catalog's password is used.

If you are not specifying new protection attributes for the alternate index and you wish to copy the model's passwords and protection attributes, you must supply the master password of either the model entry or its catalog.

catname

names the model-entry's catalog. You must identify the catalog that contains the model entry when:

- You want to specify the catalog's password instead of the model-entry's password.
- The model-entry's catalog is not identified with a JOBCAT or STEPCAT DD statement, and is not the master catalog.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved. See "Order of Catalog Use: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

ORDERED|UNORDERED

specifies whether volumes are to be used in the order in which they are listed in the VOLUMES parameter.

When you want each key range to reside on a separate volume, you can use ORDERED so that the first key range goes on the first volume, the second key range goes on the second volume, and so on.

If **ORDERED** is specified and the volumes cannot be allocated in the order specified, the command is terminated.

Abbreviations: ORD and UNORD

OWNER(ownerid)
specifies the identification of the alternate-index's owner.

ownerid
can contain 1 to 8 EBCDIC characters. When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('*CA''ROL*')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

Note to TSO users: If the owner is not identified with the OWNER parameter, the TSO user's userid becomes the ownerid.

READPW(password)
specifies a read password for the alternate index or its data or index component. The read password permits read operations against the alternate-index's data records.

When specified as a parameter of DATA or INDEX, the read password allows the user's program to open the data component or the index component for read-only processing of the component's data records (that is, the data portion of the sorted record, not the control fields VSAM inserts into stored records). For more details on how passwords can be used, see "Data Security and Protection."

password
is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

RECORDSIZE(average b maximum|4086 32600)
specifies the average and maximum length, in bytes, of an alternate-index record.

An alternate-index record can span control intervals, so RECORDSIZE can be larger than CONTROLINTERVALSIZE. average and maximum is any integer value that does not exceed the capacity of a control area, and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. You can identify the records as fixed length by specifying the same value for average and maximum.

You can use the following formulas to determine the size of the alternate-index record when the alternate index supports:

- a key-sequenced base cluster
 $RECSZ = 5 + AIXKL + (n \times BCKL)$
- an entry-sequenced base cluster
 $RECSZ = 5 + AIXKL + (n \times 4)$

where:

- RECSZ is the average record size.
- AIXKL is the alternate-key length (see the KEYS parameter).
- BCKL is the base cluster's prime-key length (you can issue the Access Method Services LISTCAT command to determine the base cluster's prime-key length).
- n = 1 when UNIQUEKEY is specified (RECSZ is also the maximum record size).
- n = the number of data records in the base cluster that contain the same alternate-key value, when NONUNIQUEKEY is specified.

When you also specify NONUNIQUEKEY, the record size you specify should be large enough to allow for as many key pointers or RBA pointers as you anticipate. The record-length values apply only to the alternate-index's data component.

If you plan to copy your alternate index using either the EXPORT or EXPORTRA command, you should be aware of the restriction in maximum record length imposed by these commands. Refer to the sections "Moving Entries" and "Copying a Catalog Entry From the Catalog Recovery Area" for details of this restriction.

Abbreviation: RECSZ

REPLICATE|NOREPLICATE

specifies whether each index record (that is, each record in the alternate-index's index component) is to be written on a direct-access device track as many times as it will fit.

When you specify REPLICATE, rotational delay is reduced and performance is improved. However, the alternate-index's index usually takes more direct-access device space.

Refer to the description of the IMBED|NOIMBED parameter for a discussion of the relationship between IMBED|NOIMBED and REPLICATE|NOREPLICATE.

Abbreviations: REPL and NREPL

REUSE|NOREUSE

specifies whether the alternate index can be used over and over again as a (temporary) new alternate index. When a reusable alternate index is opened, its high-used RBA can

be set to zero if you open it with an ACB which specifies the RESET attribute. When you use Access Method Services command, BLDINDEX, to build a reusable alternate index, the high-used RBA is always reset to zero when the alternate index is opened for BLDINDEX processing.

You cannot specify REUSE when you also specify KEYRANGES or UNIQUE for the alternate index or its components. Reusable alternate indexes may be multivolumed and are restricted to sixteen physical extents per volume.

Abbreviations: RUS and NRUS

SHAREOPTIONS(crossregion[b crosssystem][1 3])
specifies how an alternate-index's data or index component can be shared among users. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for a full description of data-set sharing.

Unless you specify sharing options, the alternate index cannot be shared.

crossregion

specifies the amount of sharing allowed among regions (within the same system). Independent job steps in an operating system can access a VSAM data set concurrently. To share a data set, each user must specify DISP=SHR in the data-set's DD statement. The values that can be specified are:

- 1 specifies that the data set can be shared by any number of users for read processing, or the data set can be accessed by only one user for read and write processing. With this option, VSAM ensures complete data integrity for the data set.
- 2 specifies that the data set can be accessed by any number of users for read processing and it can also be accessed by one user for write processing. With this option, VSAM ensures write integrity by obtaining exclusive control for a control interval when it is to be updated. If a user desires read integrity, it is his responsibility to use the ENQ and DEQ macros appropriately to provide read integrity for the data his program obtains.
- 3 specifies that the data set can be fully shared by any number of users. With this option, each user is responsible for maintaining both read and write integrity for the data his program accesses. User programs that ignore the write-integrity guidelines can cause VSAM program checks, lost or inaccessible records, uncorrectable data-set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared by any number of users, and buffers used for direct processing are refreshed for each request. This option requires your program to use the ENQ and DEQ macros to maintain data integrity while sharing the data set. Improper use of the ENQ macro can cause problems similar to those described under SHAREOPTION 3.

crosssystem

specifies the amount of sharing allowed among systems. Job steps of two or more OS/VS operating systems can

gain access to the same VSAM data set regardless of the disposition specified in each step's DD statement for the data set. To get exclusive control of the data-set's volume, a task in one system issues the RESERVE macro. The level of cross-system sharing allowed by VSAM applies only in a multiple operating system environment. The values that can be specified are:

- 1 Reserved
- 2 Reserved
- 3 specifies that the data set can be fully shared. With this option, each user is responsible for maintaining both read and write integrity for the data his program accesses. User programs that ignore write-integrity guidelines can cause VSAM program checks, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared. Buffers used for direct processing are refreshed for each request. This option requires that you use the RESERVE and RELEASE macros to maintain data integrity while sharing the data set. Writing is limited to PUT-update and PUT-insert processing that does not change the high-used RBA if your program opens the data set with DISP=SHR. Data set integrity cannot be maintained unless all jobs accessing the data set in a cross-system environment specify DISP=SHR. Improper use of the RESERVE macro can cause problems similar to those described under SHAREOPTION 3.

The assistance VSAM provides each user's program to ensure data integrity in a shared environment is:

- Each PUT request results in the appropriate buffer being written immediately into the VSAM object's direct-access device space. VSAM writes out the buffer in the user's address space that contains the new or updated data record.
- Each GET request results in all of the user's input buffers being refreshed. Each buffer's contents (that is, each data and index buffer being used by the user's program) is retrieved from the VSAM object's direct-access device.

For data integrity purposes, you should make sure that share options specified for data and index components are the same. Additional information about shared data can be found in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.

Abbreviation: SHR

SPEED|RECOVERY

specifies whether the data-component's control areas are preformatted before alternate-index records are loaded into them. When RECOVERY is specified, the data-component's control areas are written with records that indicate end-of-file. When an alternate-index record is written into a control interval, it is always followed by a record that identifies the just-written record as the last record in the alternate index. If the initial load fails, you can resume loading alternate-index records after the last

correctly-written record (because an end-of-file indicator identifies it as the last record—that is, no more alternate-index records follow).

When SPEED is specified, the data-component's space is not preformatted. Its space might contain data records from a previous use of the space, or it might contain binary zeros—its contents are unpredictable. If the initial load fails, you must load the alternate-index records again from the beginning (because VSAM is unable to determine where your last correctly-written record is—VSAM cannot find a valid end-of-file indicator when it searches your alternate-index records).

When you specify RECOVERY, your initial load takes longer because the control areas are written initially with end-of-file indicators, and again with your alternate-index records. When you specify SPEED, your initial load is quicker.

Abbreviation: RCVY

STAGE|BIND|CYLINDERFAULT

specifies how an alternate index or component that is stored on a mass storage volume is to be staged.

STAGE

indicates that the alternate index or component is to be staged from mass storage to a direct-access storage staging drive when the alternate index or component is opened. If the alternate index or component cannot be staged at open time because of heavy staging activity of other objects, data is staged as a processing program needs it.

BIND

indicates that the alternate index or component is not only to be staged, but also to be bound—that is, retained on the direct-access storage staging drive until it is closed. If the alternate index or component cannot be staged at open time because of heavy staging activity of other objects, data is staged as a processing program needs it.

CYLINDERFAULT

indicates that the alternate index or component is not to be staged when it is opened, but that data from it is to be staged as a processing program needs it.

Abbreviation: CYLF

When the alternate index or component not stored on a mass storage volume, the attribute is ineffective until the direct-access storage volume the alternate index or component is stored on is converted to a mass storage volume (by way of the CONVERTV command, which is described in OS/VS Mass Storage System (MSS) Services: Reference Information).

When one of these parameters is specified for the data component and another parameter is specified for the index component, the components are staged separately as specified, except when the sequence set of the index component is imbedded in the data. In that case, the parameter specified for the index component applies to both components.

TO(date)|FOR(days)

specifies the retention period for the alternate index. The alternate index is not automatically deleted when the expiration date is reached. When you do not specify a retention period, the alternate index can be deleted at any time.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the Julian date (001, for January 1, through 365, for December 31), through which the alternate index is to be kept before it is allowed to be deleted.

FOR(days)

specifies the number of days for which the alternate index is to be kept before it is allowed to be deleted. days can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. The maximum number that can be specified for days is 9999. If the number specified is 0 through 1830, the retention period is the number of days specified. If the number specified is between 1831 and 9999, the retention period is through the year 1999.

UNIQUE|SUBALLOCATION

specifies whether the alternate-index's components are allocated an amount of space from the volume's available space (UNIQUE) or from a VSAM data-space's available space (SUBALLOCATION.)

When UNIQUE is specified, a VSAM data space is built and is assigned exclusively to each component of the alternate index. The data space is created when the alternate index is defined. The alternate-index's volume(s) must be mounted. VSAM builds a DSCB in the volume's table of contents (VTOC) to describe the data space. The name of the data space, which is the same as the component's name, is put in the DSCB. A subentry is added to the volume entry (in the VSAM catalog) to describe the VSAM data space.

When SUBALLOCATION is specified, space from one of the VSAM data spaces on the volume is assigned to the alternate-index's component.

The space-allocation attribute interacts with other DEFINE ALTERNATEINDEX parameters. You should take care to ensure that the attribute you specify for the alternate index is consistent with other attributes:

- REUSE: You cannot specify REUSE when you specify UNIQUE for an alternate index or its components.
- KEYRANGES: When UNIQUE is specified, a data space is built and allocated on a separate volume for each key range.
- VOLUMES: When UNIQUE is not specified, VSAM data space must exist on the volume that is to contain the alternate-index's component. When UNIQUE is specified, and more than one volume is specified, VSAM must already own all the volumes except the first. If there is no VSAM space on a volume, you must execute a DEFINE SPACE CANDIDATE before your DEFINE UNIQUE.

Abbreviations: UNQ and SUBAL

UNIQUEKEY|NONUNIQUEKEY

specifies whether more than one data record (in the base cluster) can contain the same alternate-key value.

UNIQUEKEY specifies that each alternate key points to only one data record. When the alternate index is built (see the BLDINDEX command) and more than one data record contains the same alternate-key value, the BLDINDEX processing terminates with an error message.

NONUNIQUEKEY specifies that an alternate-key value might point to more than one data record in the base cluster. The alternate-key record can point to a maximum of 32,768 records with nonunique keys.

When you specify NONUNIQUEKEY, the value you specify as the maximum record size should be large enough to allow for alternate-index records that point to more than one data record.

Abbreviations: UNQK and NUNQK

UPDATEPW(password)

specifies the update password for the alternate index or its data or index component. The update password permits read and write operations against the alternate-index's data records.

When specified as a parameter of DATA or INDEX, the update password allows the user's program to open the data component or index component for read and write processing of the component's VSAM records (that is, the data portion of the stored record, not the control fields VSAM inserts into stored records).

If a read password is the only password specified for the object (that is, it is the highest-level password), it propagates upward and becomes the password for all higher levels. If you specify a higher-level password and do not specify an update password, the update or read password is null.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, UPDATEPW('XDORISX')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, UPDATEPW('XCA''ROLX')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, UPDATEPW(X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

UPGRADE|NOUPGRADE

specifies whether the alternate-index is to be upgraded (that is, kept up to date) when its base cluster is modified. UPGRADE specifies that when the base-cluster's records are added to, updated, or erased, the cluster's alternate index is upgraded to reflect the changed data.

When UPGRADE is specified, the alternate index's name is cataloged with the names of other alternate indexes for the base cluster. The group of alternate-index names identify the upgrade set—all of the base cluster's alternate indexes that are opened when the base cluster is opened for write operations.

The UPGRADE attribute is not effective for the alternate index until the alternate index is built (see the BLDINDEX command). If the alternate index is defined when the base cluster is open, the UPGRADE attribute takes effect the next time the base cluster is opened.

Abbreviations: UPG and NUPG

WRITECHECK | NOWRITECHECK

specifies whether an alternate-index's data or index component is to be checked (by a direct-access device operation called write-check) when a record is written to the device. If WRITECHECK is specified, the record is written to the device, then read without data transfer to test for the data-check condition.

Abbreviations: WCK and NWCK

Components (Alternate Index)

Attributes can be specified separately for the alternate-index's data and index components.

DATA(options)

specifies attributes for the alternate-index's data component. Except for name, space allocation, and protection attributes, attributes specified as parameters of DATA override the same attributes specified as parameters of ALTERNATEINDEX. Attributes specified as parameters of DATA can also be used to override the same attributes of a data entry that is used as a model.

Parameters that can be specified for the alternate-index's data component are:

- **ATTEMPTS, AUTHORIZATION, and CODE:** Specifies protection attributes for the data component when it is opened as a separate VSAM object by the user's program.
- **BUFFERSPACE:** Specifies the amount of buffer space that the user's program is to provide from its address space when the program opens the data component as a separate VSAM object.
- **CONTROLINTERVALSIZE:** Specifies the size of the data component's control interval. The sizes you can specify are:

$CISZ = (n \times 512) \text{ or } (n \times 2048)$

where n is a positive integer from 1 to 16.

- **CYLINDERS, RECORDS, or TRACKS:** Specifies an amount of space to be allocated to the data component. When you specify space as a parameter of DATA, you cannot also specify space for the alternate index as a whole. (See the description of TRACKS | CYLINDERS | RECORDS as parameters of ALTERNATEINDEX.)
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the data component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that closes it.
- **ERASE or NOERASE:** Specifies whether the data records are to be erased when the alternate index is deleted.
- **EXCEPTIONEXIT:** Identifies an exception exit routine for the alternate index's data component.
- **FILE:** Names the DD statement that describes a direct-access device and volumes that are to contain the data component. When more than one volume is to contain the data component, all volumes must be of the same device type.
- **FREESPACE:** Specifies the amount of space that is to be left free after any allocation and after any split of control intervals and control areas when the alternate

index is built (see the BLDINDEX command).

- **KEYRANGES:** Specifies that portions, or key ranges, of the alternate index are to be put on different volumes, and specifies the boundaries of each key range.
- **KEYS:** Describes the alternate-key field of the base cluster's data record.
- **MASTERPW, CONTROLPW, UPDATEPW, or READPW:** Specifies passwords for the data component when it is opened as a separate VSAM object by the user's program.
- **MODEL:** Identifies a data entry that is to be used as a model for the data component only.
- **NAME:** Names the data component. If you do not specify a name for the data component, VSAM generates one (see "Defining a Cluster" for the format of the generated name).
- **ORDERED or UNORDERED:** When ORDERED is specified, the volumes listed with the VOLUMES parameter are allocated in the order listed.
- **OWNER:** Identifies the data component's owner, when the owner is different from the owner of the alternate index as a whole.
- **RECORDSIZE:** Specifies the average and maximum sizes of the alternate-index's data records.
- **REUSE or NOREUSE:** Specifies whether the alternate index's data component is to be reusable.
- **SHAREOPTIONS:** Specifies share options for the data component when it is opened as a separate VSAM object by the user's program.
- **STAGE, BIND, or CYLINDERFAULT:** Specifies whether the data component, stored on a mass storage volume, is to be staged.
- **UNIQUE or SUBALLOCATION:** Specifies whether the data component is to occupy exclusively a VSAM data space or share an existing data space with other VSAM objects.
- **UNIQUEKEY or NONUNIQUEKEY:** Specifies whether more than one data record (in the base cluster) can contain the same alternate key value.
- **VOLUMES:** Identifies the volume(s) that is to contain the data component, and identifies other volumes that can be used as overflow volumes for the data component.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the data records are to be checked by the direct-access device when they are written into the data component.

Restrictions are noted with each parameter's description. All of the DATA subparameters are described in more detail as parameters of the alternate index as a whole.

INDEX(options)

specifies attributes for the index component of an alternate index.

Except for name, space allocation, and protection attributes, attributes specified as parameters of INDEX override the same attributes specified as parameters of ALTERNATEINDEX. Attributes specified as parameters of INDEX can also be used to override the same attributes of an index entry that is used as a model.

Parameters that can be specified for the alternate-index's index component are:

- **ATTEMPTS, AUTHORIZATION, and CODE:** Specifies protection attributes for the index when it is opened as a separate VSAM object by the user's program.
- **CONTROLINTERVALSIZE:** Specifies the index component's control interval size. The sizes you can specify are: 512, 1024, 2048, or 4096 bytes.
- **CYLINDERS, RECORDS, or TRACKS:** Specifies an amount of space to be allocated to the index component. When you specify one of the space-quantity parameters as a parameter of INDEX, you must also specify a space-quantity parameter as a parameter of DATA. You cannot also specify a quantity of space for the alternate index as a whole.
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the index component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that closes it.
- **EXCEPTIONEXIT:** Identifies an exception exit routine for the index component.
- **FILE:** The DD statement describes a direct-access device that is to contain the index component. When more than one volume is to contain the index component, all volumes must be of the same device type.
- **IMBED or NOIMBED:** Specifies whether the sequence set records are to be written adjacent to the data.
- **MASTERPW, CONTROLPW, UPDATEPW, or READPW:** Specifies passwords for the index component when it is opened as a separate VSAM object by the user's program.
- **MODEL:** Identifies an index entry that is to be used as a model for the index component only.
- **NAME:** Names the index component. If you do not specify a name for the index component, VSAM generates one (see "Defining a Cluster" for the format of the generated name).
- **ORDERED or UNORDERED:** When ORDERED is specified, the first volume in the volser list (of the VOLUMES parameter) is to contain the index. Other volumes are to be available as overflow volumes. Otherwise, any volume specified in the volser list can be used for the index component.
- **OWNER:** Identifies the index component's owner, when the owner identification is different from the owner of the alternate index as a whole.
- **REPLICATE or NOREPLICATE:** Specifies whether the index component's records are to be repeated around the track.
- **REUSE or NOREUSE:** Specifies whether the alternate index's index component is to be reusable.
- **SHAREOPTIONS:** Specifies share options for the index component when it is opened as a separate VSAM object by the user's program.
- **STAGE, BIND, or CYLINDERFAULT:** Specifies how the index component, stored on a mass storage volume, is to be staged.

- **UNIQUE or SUBALLOCATION:** Specifies whether the index component is to occupy exclusively a VSAM data space or share an existing data space with other VSAM objects.
- **VOLUMES:** Identifies the volume that is to contain the index component, and identifies other volumes that can be used as overflow volumes for the index component.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the index component's records are to be checked by the direct-access device when they are written into the index.

Restrictions are noted above with each parameter's description. All of the INDEX subparameters are described in more detail as parameters of the alternate index as a whole.

Catalog (Alternate Index)

The catalog parameter allows you to supply the name and password of the catalog in which the alternate index is to be defined. When you specify CATALOG, you identify the catalog that contains the base-cluster's entry.

CATALOG(catname[/password])

identifies the catalog in which the alternate index is to be defined. The catalog also contains the base cluster's entry (see the RELATE parameter above). See "Order of Catalog Use: DEFINE" for information about the order in which a catalog is selected if the catalog's name is not specified.

catname

specifies the catalog's name.

password

specifies the catalog's update or higher level password. If the catalog is password-protected, you must supply the catalog's update or higher level password. If no password is specified with the CATALOG parameter, VSAM asks the operator for the catalog's update password.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved.

If the base cluster's master password is not specified with the RELATE parameter, then the catalog's master password must be specified.

Abbreviation: CAT

DEFINE ALTERNATEINDEX EXAMPLE

Define an Alternate Index: Example 1

In this example, an alternate index is defined. An example for DEFINE CLUSTER illustrates the definition of the alternate-index's base cluster, EXAMPLE.KSDS2. A subsequent example illustrates the definition of a path, EXAMPLE.PATH, that allows you to process the base-cluster's data records using the alternate key to locate them. The alternate index, path, and base cluster are defined in the same catalog, AMASTCAT.

```
//DEFAIX JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
        DEFINE ALTERNATEINDEX -
            (NAME(EXAMPLE.AIX) -
             RELATE(EXAMPLE.KSDS2) -
             MASTERPW(AIXMRPW) -
             UPDATEPW(AIXUPPW) -
             KEYS(3 0) -
             RECORDSIZE(40 50) -
             VOLUMES(VSER04) -
             CYLINDERS(3 1) -
             NONUNIQUEKEY -
             UPGRADE) -
            CATALOG(AMASTCAT/MCATUPPW)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE ALTERNATEINDEX command creates an alternate-index entry, a data entry, and an index entry to define the alternate index EXAMPLE.AIX. The DEFINE ALTERNATEINDEX command also obtains space for the alternate index from one of the VSAM data spaces on volume VSER04, and allocates three cylinders for the alternate index's use. Since the alternate index is being defined into a recoverable catalog, the catalog recovery volume will be dynamically allocated. The command's parameters are:

- NAME, which specifies that the alternate-index's name is EXAMPLE.AIX.
- RELATE, which identifies the alternate-index's base cluster, EXAMPLE.KSDS2.
- MASTERPW and UPDATEPW, which specifies the alternate-index's master password, AIXMRPW, and update password, AIXUPPW.
- KEYS, which specifies the length and location of the alternate key field in each of the base-cluster's data records. the alternate key field is the first three bytes of each data record.
- RECORDSIZE, which specifies that the alternate-index's records are variable-length, with an average size of 40 bytes and a maximum size of 50 bytes.
- VOLUMES, which specifies that the alternate index is to reside on volume VSER04. This example assumes that the volume is already cataloged in the master catalog, AMASTCAT.
- CYLINDERS, which specifies that three cylinders are allocated for the alternate-index's space. When the alternate index is extended, it is to be extended in increments of one cylinder.

- NONUNIQUEKEY, which specifies that the alternate key value might be the same for two or more data records in the base cluster.
- UPGRADE, which specifies that the alternate index is to be opened by VSAM and upgraded each time the base cluster is opened for processing.
- CATALOG, which specifies that the alternate index is to be defined in the master catalog, AMASTCAT. The example also supplies the master catalog's update password, MCATUPPW.

DEFINE CLUSTER

When you define a cluster, you can specify attributes for the cluster as a whole and for the cluster's components. The general format of the DEFINE CLUSTER command is:

DEFINE CLUSTER (parameters) -

[DATA(parameters)] -

[INDEX(parameters)] -

[CATALOG(subparameters)]

DEFINE CLUSTER

DEFINE	CLUSTER
	(NAME(entryname) {CYLINDERS(primary[b secondary])} RECORDS(primary[b secondary])} TRACKS(primary[b secondary])} VOLUMES(volser[b volser...]) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [DESTAGWAIT <u>NODESTAGWAIT</u>] [ERASE <u>NOERASE</u>] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE(CI-percent [b CA-percent] 0 0)] [IMBED <u>NOIMBED</u>] [INDEXED <u>NONINDEXED</u> NUMBERED] [KEYRANGES({lowkey b highkey [(lowkey b highkey)...]})] [KEYS(length b offset 64 0)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [ORDERED <u>UNORDERED</u>] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE(average b maximum)] [REPLICATE <u>NOREPLICATE</u>] [REUSE <u>NOREUSE</u>] [SHAREOPTIONS(crossregion [b crosssystem] 1 3)] [SPANNED <u>NONSPANNED</u>] [SPEED <u>RECOVERY</u>] [STAGE BIND CYLINDERFAULT] [TO(date) FOR(days)] [UNIQUE <u>SUBALLOCATION</u>] [UPDATEPW(password)] [WRITECHECK <u>NOWRITECHECK</u>])

(Continued on next page.)

	<pre> [DATA [[ATTEMPTS(number)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [CYLINDERS(primary[b secondary]] RECORDS(primary[b secondary]] TRACKS(primary[b secondary]])] [DESTAGWAIT NODESTAGWAIT] [ERASE NOERASE] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE(CI-percent[b CA-percent])] [KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])] [KEYS(length b offset)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [NAME(entryname)] [ORDERED UNORDERED] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE(average b maximum)] [REUSE NOREUSE] [SHAREOPTIONS(crossregion[b crosssystem])] [SPANNED NONSPANNED] [SPEED RECOVERY] [STAGE BIND CYLINDERFAULT] [UNIQUE SUBALLOCATION] [UPDATEPW(password)] [VOLUMES(volser[b volser...])] [WRITECHECK NOWRITECHECK]]] [INDEX [[ATTEMPTS(number)] [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [CYLINDERS(primary[b secondary]] RECORDS(primary[b secondary]] TRACKS(primary[b secondary]])] [DESTAGWAIT NODESTAGWAIT] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [IMBED NOIMBED] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [NAME(entryname)] [ORDERED UNORDERED] [OWNER(ownerid)] [READPW(password)] [REPLICATE NOREPLICATE] [REUSE NOREUSE] [SHAREOPTIONS(crossregion[b crosssystem])] [STAGE BIND CYLINDERFAULT] [UNIQUE SUBALLOCATION] [UPDATEPW(password)] [VOLUMES(volser[b volser...])] [WRITECHECK NOWRITECHECK]]] [CATALOG(catname[/password])] </pre>
--	---

DEFINE CLUSTER Parameters: Summary

The parameters of the DEFINE CLUSTER command can be divided into required and optional parameters.

The parameters which are required for all types of clusters are:

- Entry type, which specifies that a cluster and its catalog entry are to be created (CLUSTER).
- Name, which names the cluster (NAME).
- Allocation, which specifies the amount of space to be allocated to the cluster and identifies the volume(s) that is to contain the cluster (VOLUMES and one of CYLINDERS, RECORDS, or TRACKS). You need not explicitly specify the allocation parameter if you specify the MODEL parameter.

The parameters which are required for an entry-sequenced or relative-record cluster but are optional for a key-sequenced cluster are:

- Type, which specifies the kind of data organization for the cluster (INDEXED, NONINDEXED, or NUMBERED). You need not explicitly specify the type parameter for a key-sequenced cluster. You need not explicitly specify the type parameter for an entry-sequenced or relative-record cluster if you specify the MODEL parameter.

The optional parameters are:

- Model, which identifies an existing cluster entry that is to be used as a model for the cluster being defined (MODEL).
- Allocation, which specifies
 - whether the cluster's data space is to be shared with other VSAM objects (SUBALLOCATION or UNIQUE).
 - the name of the DD statement that describes the cluster's volume(s) (FILE).
 - the size of control intervals, records, and the user program's buffer area (CONTROLINTERVALSIZE, RECORDSIZE, and BUFFERSPACE).
 - how a cluster or component that is stored on a mass storage device is to be staged (STAGE, BIND, or CYLINDERFAULT).
 - whether the volume's are to be allocated in the order given with the VOLUMES parameter (ORDERED or UNORDERED).
 - whether the cluster is intended for temporary storage of data (REUSE or NOREUSE).
 - whether the cluster's records can cross control-interval boundaries (SPANNED or NONSPANNED).
- Data integrity, which specifies
 - the name of a user-written I/O error routine (EXCEPTIONEXIT).
 - whether data records are to be overwritten with binary zeros when the cluster is deleted (ERASE or NOERASE).
 - whether the data records are to be checked when the direct-access device writes them into secondary storage (WRITECHECK or NOWRITECHECK).
 - whether a cluster or component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that

closes it (DESTAGEWAIT or NODESTAGEWAIT).

- whether space is to be preformatted when the data is initially loaded (RECOVERY or SPEED).
- Protection, which
 - specifies passwords for the cluster (MASTERPW, CONTROLPW, UPDATEPW, and READPW).
 - identifies the cluster to a console operator without revealing the cluster's entryname (ATTEMPTS and CODE).
 - identifies a user-written-security-verification routine for additional authorization verification (AUTHORIZATION).
 - identifies the cluster's owner (OWNER).
 - specifies a retention period for the cluster (FOR or TO).
- Parameters that apply only to the definition of a key-sequenced cluster, which specify
 - whether index records are to be repeated around the track of the index's direct-access device (REPLICATE or NOREPLICATE).
 - whether the index's sequence-set records are to be written adjacent to their data control areas (IMBED or NOIMBED).
 - whether the data records are to be divided by key among volumes (KEYRANGES).
 - the length and offset of the cluster's key field (KEYS).
 - the amount of free space to be left in control intervals and control areas when the data records are loaded (FREESPACE).
- Components, which specifies attributes for the cluster's data and, for a key-sequenced cluster, index components that might differ from those attributes specified for the cluster as a whole or the attributes established for the cluster being used as a model (DATA and INDEX).
- Catalog, which identifies the catalog in which the cluster is to be defined (CATALOG).

DEFINE can be abbreviated: DEF.

DEFINE CLUSTER PARAMETERS

Required Parameters

CLUSTER

specifies that a cluster is to be defined. CLUSTER is followed by the parameters specified for the cluster as a whole; these parameters are enclosed in parentheses and, optionally, are followed by the DATA and/or INDEX parameters and their subparameters, enclosed in parentheses.

Abbreviation: CL

NAME(entryname)

specifies the name of the entry (cluster or component) being defined. NAME must be specified for the cluster; it can optionally be specified for a data or index component. If no name is specified for the data or index component, a

name is generated. Because the cluster, data component, and index component are individually named, each can be addressed.

The name may contain from 1 through 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and the 12-0 overpunch). Names containing more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of any name or name segment must be either an alphabetic or a national character. If you specify the name in the format `entryname(modifier)`, only that portion of the name preceding the left parenthesis will be used. The modifier enclosed in parentheses will be ignored.

`CYLINDERS(primary[b secondary])`
`RECORDS(primary[b secondary])`
`TRACKS(primary[b secondary])`

specifies the amount of space to be allocated to the cluster from the volume's available space, when `UNIQUE` is specified, or from the available space of one of the volume's VSAM data spaces. If you specify `RECORDS`, the amount of space allocated is the minimum number of tracks that are required to contain the specified number of records. However, if you specify `TRACKS` or `RECORDS` and if the minimum number of tracks should exceed a cylinder, space is allocated in terms of cylinders.

Define by records will allocate sufficient space to contain that number of records, but variable factors unknown at define time (such as key compression or method of loading records) may result in inefficient use of the space allocated, preventing every data CA from being completely used, and causing the user to be unable to load that number of records without requiring secondary allocation. When `UNIQUE` is specified and the data component is to be contained on more than one volume, each volume can contain a maximum of 16 extents.

When the data component is not divided into key ranges and more than one volume is specified, the primary amount of space is allocated only on the first volume when the component is defined. The secondary amount must be coded to extend the component to additional volumes. When the component increases in size so as to extend to additional volumes, the first allocation on each overflow volume is the primary amount.

When `KEYRANGES` is specified and the data component is to be contained on more than one volume, the primary amount of space is immediately allocated on each volume required for the key ranges. If more than one key range is allocated on a single volume, the primary amount is allocated for each key range.

secondary amounts can be allocated on all volumes available to contain parts of the cluster regardless of the key ranges. The primary and each secondary allocation must be able to be satisfied in 5 extents; otherwise your `DEFINE` or data set extension will fail.

If you do not specify the `MODEL` parameters, you must specify one, and only one, of the following parameters: `CYLINDERS`, `RECORDS`, or `TRACKS`. You can specify the amount of space as a parameter of `CLUSTER`, as a parameter of `DATA`, or as a parameter of both `DATA` and `INDEX`. When a key-sequenced cluster is being defined and the space is specified as a parameter of:

- `CLUSTER`, the amount specified is divided between the data and index components. The division algorithm is a function of control interval size, record size, device type, and other data set attributes. If the division

results in an allocation for the data component which is not an integral multiple of the required control-area size, the data-component's allocation is rounded up to the next higher control-area multiple. This rounding may result in a larger total allocation for your cluster than which you specified.

- DATA, the entire amount specified is allocated to the data component. An additional amount of space, depending on control-interval size, record size, device type, and other data set attributes, is allocated to the index component.

When you specify UNIQUE, and the amount of space you specify for the data component of the cluster is not an even multiple of cylinder-size, VSAM rounds the allocation upward to a multiple of a cylinder.

To determine the exact amount of space allocated to each component, list the cluster's catalog entry, using the LISTCAT command.

When you specify UNIQUE, and the cluster's data space is the first data space on a volume that belongs to a recoverable catalog, an additional amount of (the equivalent, in tracks, of) one cylinder is allocated for the recovery area data space.

primary

specifies the initial amount of space that is to be allocated to the cluster.

secondary

specifies the amount of space that is to be allocated each time the cluster extends, as a secondary extent. When secondary is specified, space for the cluster's data and index components can be expanded to include a maximum of 123 extents. If the data space that contains (all or part of) the cluster cannot be extended because the cluster's secondary allocation amount is greater than the data-space's secondary allocation amount, VSAM builds a new data space. The new data-space's primary and secondary allocation amounts are equal to the cluster's secondary allocation amount.

primary and secondary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviations: CYL, REC, and TRK

VOLUMES(volser[b volser...])

specifies the volumes to contain the cluster or component. VOLUMES must be specified if you do not specify the MODEL parameter. It can be specified for the cluster as a whole. Alternatively, it can be specified as a parameter of DATA and, if the cluster is key-sequenced, as a parameter of INDEX.

If the data and index components are to reside on different device types, VOLUMES must be specified as both a DATA and an INDEX parameter. If more than one volume is listed with a single VOLUMES parameter, the volumes must be of the same device type. In a system with the Mass Storage System, a cluster or component can be defined on a mass storage volume.

A volume serial number may be repeated in the list only if the INDEXED and KEYRANGE parameters are specified. You may wish to do this in order to have more than one key range on the same volume. Even in this case, repetition is only valid if all duplicate occurrences are used for the primary allocation of some key range.

volser

is a 1 to 6 alphanumeric or national character volume serial number. Special characters you can specify are the comma [,], semicolon [;], parentheses [(and)], slash [/], asterisk [*], period [.] , single quotation mark ['], ampersand [&], plus sign [+], hyphen [-], and equal sign [=]. Code a single quotation mark as two single quotation marks.

If the volume serial number includes a special character, enclose volser in quotation marks. For example, you can specify VOLUMES('VOL×10' 'VOL'AB') to include volumes VOL×10 and VOL'AB in the cluster's volume list.

The VOLUMES parameter interacts with other DEFINE CLUSTER parameters. You should take care to ensure that the volume(s) you specify for the cluster can satisfy the cluster's other attributes:

- SUBALLOCATION: If UNIQUE is not specified, the volume must contain VSAM data space.
- KEYRANGES: If KEYRANGES and UNIQUE are specified, a VSAM data space is built and allocated for each key range on a separate volume.
- ORDERED: If ORDERED is specified, the volumes are used in the order listed. If ORDERED and KEYRANGES are specified, there is a one-for-one correspondence between the key ranges in the key-range list and the volumes in the volser list.
- CYLINDERS, RECORDS, TRACKS: The volume(s) must contain enough unallocated space to satisfy the component's primary space requirement.
- FILE: The volume information supplied with the DD statement(s) pointed to by FILE must be consistent with the information specified for the cluster and its components.
- CATALOG: The data space on the volume must have been defined in the same catalog that the cluster is being defined in, and must be owned by the catalog.

Abbreviation: VOL

Data Organization

When a cluster is defined, you specify whether the data is to be indexed (key-sequenced), nonindexed (entry-sequenced), or numbered (relative-record).

Certain parameters apply only to key-sequenced clusters, as noted in the description of each of these parameters.

If you do not specify the data organization and you also do not specify the MODEL parameter, your cluster will default to key-sequenced (indexed). If you wish to define either an entry-sequenced or relative-record cluster, you must specify the NONINDEXED or NUMBERED parameter unless you specify the MODEL parameter.

INDEXED|NONINDEXED|NUMBERED

specifies the type of data organization that the cluster is to have. When a cluster is defined, you can specify whether the data organization is to be indexed (key-sequenced), sequential (entry-sequenced), or numbered (relative-record). The data organization you select must be consistent with other parameters you specify.

INDEXED

specifies that the cluster being defined is for key-sequenced data. If INDEXED is specified, an index component is automatically defined and cataloged. The data records can be accessed by key or by relative-byte address (RBA). See also the INDEX parameter.

Abbreviation: IXD

NONINDEXED

specifies that the cluster being defined is for entry-sequenced data. The data records can be accessed sequentially or by relative-byte address (RBA).

Abbreviation: NIXD

NUMBERED

specifies that the cluster's data organization is for relative-record data. A relative-record-cluster is similar to an entry-sequenced cluster, and has fixed-length records that are stored in slots. Each slot might or might not contain a data record. The data records are accessed by relative-record number (that is, slot number).

Abbreviation: NUMD

Optional Parameters**ATTEMPTS(number|2)**

specifies the maximum number of times the operator can try to enter a correct password in response to a prompting message. The number can be any number, 0 through 7. If 0 is specified, the operator is not prompted and is not allowed to enter a password from the console.

number

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. This parameter can be coded, but only has effect when the entry's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

Note to TSO Users: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password for the cluster. Checking the logon password counts as one attempt to obtain a password. If ATTEMPTS is not specified, the user has one attempt to supply the cluster's password because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])

specifies that a user-security-verification routine (USVR) is available for additional security verification. When a protected cluster is accessed and the user supplies a correct password other than the cluster's master password, the USVR receives control. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for details on the USVR.

If a USVR is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

This parameter can be coded, but only has effect when the entry's master password is not null.

entrypoint

specifies the name of the user's-security-verification routine.

The name can contain 1 to 8 alphameric characters (A through Z, or 0 through 9), national characters (à, #, and \$), or special characters (the hyphen and 12-0 overpunch). The first character must be an alphabetic or national character.

string

specifies information to be passed to the user-security-verification routine when it receives control to verify authorization, and can contain 1 to 255 EBCDIC characters. When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single-quotation marks (for example, AUTHORIZATION(entrypoint, '*DORIS*')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, '*CA''ROL*')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint, SUE). The string can contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

BUFFERSPACE(size)

specifies the minimum space to be provided for buffers. The buffer space size you specify helps VSAM determine the data-component's and index-component's control-interval size. If BUFFERSPACE is not coded, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key sequenced, one index-component control interval.

size

is the amount of space to be provided for buffers. size can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. The size specified cannot be less than enough space to contain two data-component control intervals and, if the data is key sequenced, one index control interval. If the specified size is less than VSAM requires for the buffers needed to run your job, VSAM terminates your DEFINE and provides an appropriate error message.

Abbreviations: BUFSP or BUFSPC

CODE(code)

specifies a code name for the entry being defined. If an attempt is made to access a password-protected entry without a password, the code name is used in a prompting message; the code enables the operator to be prompted for the password without disclosing the name of the entry.

If CODE is not specified and an attempt is made to access a cluster or component that is password protected without supplying a password, the operator is prompted with the name of the entry.

This parameter can be coded, but only has effect when the cluster's or component's master password is not null. A prompting message is issued only when the user hasn't already supplied the appropriate password.

code

can contain 1 to 8 alphanumeric or special characters. When the code contains a special character (that is, a comma [,], semicolon [;], blank [], period [.], single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the code in single-quotation marks (for example, CODE(' *DORIS *')).

If the code contains a lowercase alphabetic character it will be folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE(' *CA ' 'ROL *')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, CODE(X'E2E4C5') is the same as CODE(SUE).

CONTROLINTERVALSIZE(size)

specifies the size of the control interval for the cluster or component.

The size of the control interval depends on the maximum size of the data records and the amount of buffer space you provide. If the size you specify is not an integer multiple of 512 or 2048, VSAM selects the next higher multiple for your cluster's control-interval size.

If CONTROLINTERVALSIZE is not coded, VSAM determines the size of control intervals. If you have not specified BUFFERSPACE and the size of your records permits, VSAM uses the optimum size for the data component and 512 for the index component.

size

for a cluster's data component.

The size of a data-control interval must be at least 7 bytes larger than the maximum record length if the data set does not have the SPANNED attribute. If the control interval specified is less than maximum record length plus the 7-byte overhead, then VSAM will increase the data-control interval size to contain the maximum record length plus the needed overhead. If the data set has the spanned attribute, the control-interval size can be less than the maximum record length. The sizes you can specify—between 512 and 32,768 bytes—are an integer multiple of 512 or 2048:

$$CISZ = (n \times 512) \text{ or } (n \times 2048)$$

where n is a positive integer from 1 to 16

size

for a key-sequenced cluster's index component.

You can specify the following values:

$$CISZ = [512 \mid 1024 \mid 2048 \mid 4096]$$

When you specify a size that is not a multiple of 512 or 2048, VSAM chooses the next higher multiple.

Refer to the chapter "Optimizing VSAM's Performance" in OS/VS Virtual Storage Access Method (VSAM) Programmer's

Guide for a discussion of the relationship between control interval size and physical block size. The discussion also includes restrictions that apply to control interval size and physical block size.

Abbreviations: CISZ or CNVSZ

CONTROLPW(password)

specifies a control password for the entry being defined. The control password permits operations using control-interval access to read, write, and update control intervals that contain the cluster's records.

If a read or update password is the only password specified for the object, it (the highest-level password) propagates upward and becomes the password for all higher unspecified levels.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

DESTAGEWAIT | NODESTAGEWAIT

specifies whether a cluster or its component that is stored on a mass storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it.

DESTAGEWAIT indicates that destaging is to be completed before VSAM returns control to the program that issued the CLOSE macro. VSAM can thus notify the program whether destaging was successful.

NODESTAGEWAIT indicates that notification of unsuccessful destaging is to be made only by a message to the operator.

When the cluster or component is not stored on a mass storage volume, the attribute is ineffective until the direct-access storage volume the cluster or component is stored on is converted to a mass storage volume (by way of the CONVERTV command, which is described in OS/VS Mass Storage System (MSS) Services: Reference Information).

When one of these parameters is specified for the data component and the other parameter is specified for the index component, the components are destaged separately as specified, except when the sequence set of the index component is imbedded in the data. In that case, the parameter specified for the index component applies to both components.

Abbreviations: DSTGW and NDSTGW

ERASE|NOERASE

specifies whether the cluster's data component is to be erased when its entry in the catalog is deleted. If ERASE is specified, the data component is overwritten with binary zeros when its catalog entry is deleted.

Abbreviations: ERAS and NERAS

EXCEPTIONEXIT(entrypoint)

specifies the name of the user-written routine, called the exception exit routine, that receives control when an exceptional I/O error condition occurs during the transfer of data between your program's address space and the cluster's direct-access storage space. (An exception is any I/O error condition that causes a SYNAD exit to be taken.) The component's exception-exit routine is processed first, then the user's SYNAD-exit routine receives control.

If an exception exit routine is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

Abbreviation: EEXT

FILE(dname)

specifies the name of the DD statement that identifies the devices and volumes to be used for space allocation. When you specify UNIQUE, FILE can identify the DD statement which identifies and allocates the volumes of the cluster. If more than one volume is specified, all volumes must be the same device type. If data and index components are to reside on different device types, FILE is specified as a parameter of DATA and INDEX so that separate DD statements can be referenced.

When the cluster is defined in a recoverable catalog, and FILE is specified as a parameter of CLUSTER, the FILE parameter can identify the DD statement for all volumes on which space is to be allocated. When FILE refers to more than one volume of the same device type, the DD statement that describes the volumes cannot be a concatenated DD statement. When the cluster is cataloged in a recoverable catalog, the DD statement can be concatenated if the volumes are of different device types: part of the concatenated DD statement describes the cluster's or component's volume(s) of one device type; the other part of the DD statement describes the volume(s) of another device type.

If the FILE parameter is not specified, the required volumes are dynamically allocated. If the cluster or one of its components is unique, and if the FILE parameter is not specified and the volume(s) is physically mounted, the volume(s) identified with the VOLUMES parameter is dynamically allocated. The volume(s) must be mounted as permanently resident or reserved.

The DD statement you specify must be in the form:

```
//ddname DD UNIT=(devtype[,unitcount]),
//      VOL=SER=(volser1,volser2,volser,...),...
```

FREESPACE(CI-percent[b CA-percent]|0 0)

specifies the amount of space that is to be left free when the cluster is loaded and after any split of control intervals (CI-percent) and control areas (CA-percent). The amount of empty space in the control interval and control area is available for data records that are updated and inserted after the cluster is initially loaded. This parameter applies only to key-sequenced clusters.

The amounts are specified as percentages. CI-percent translates into a number of bytes that is equal to, or slightly less than, the percentage value of CI-percent. CA-percent translates into a number of control intervals that is equal to, or less than, the percentage value of CA-percent. CI-percent and CA-percent, which must be equal to or less than 100, can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

When you specify FREESPACE(100 100), one data record is placed in each control interval used for data and one control interval in each control area is used for data (that is, one data record is stored in each control area when the data set is loaded). When no FREESPACE value is coded, the default specifies that no free space is to be reserved when the data set is loaded.

When you define the cluster using the RECORDS parameter, the amount of free space specified is not taken into consideration in the calculations to determine primary allocation.

Abbreviation: FSPC

IMBED|NOIMBED

specifies whether the sequence set (the lowest level of the index) is to be placed with the data component. If IMBED is coded, each sequence-set record for each control area is written as many times as it will fit on the first track adjacent to the control area. This parameter applies only to key-sequenced clusters.

The IMBED|NOIMBED parameter interacts with the REPLICATE|NOREPLICATE parameter in determining the physical attributes of the index of the cluster. These index attributes can be used in the following combinations:

- The sequence set records are adjacent to the data control areas, and only the sequence set records are replicated (IMBED and NOREPLICATE).
- The sequence set records are adjacent to the data control intervals, and all levels of index records are replicated (IMBED and REPLICATE).
- All index records are together, and all index records are replicated (REPLICATE and NOIMBED).
- All index records are together, and no index records are replicated (NOREPLICATE and NOIMBED).

For some applications, specifying index options can improve the application's performance.

See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for information on how the index's optional attributes affect performance.

Abbreviations: IMBD and NIMBD

**KEYRANGES((lowkey b highkey)
[[lowkey b highkey]...])**

specifies that portions of key-sequenced data are to be placed on different volumes. Each portion of the data is called a key range. This parameter applies only to key-sequenced clusters.

The maximum number of key ranges is 123. Key ranges must be in ascending order, and are not allowed to overlap. A gap can exist between two key ranges, but you are not allowed to insert records whose keys are within the gap.

The KEYRANGES parameter interacts with other DEFINE CLUSTER parameters. You should take care to ensure that, when you specify KEYRANGES, the cluster's other attributes can be satisfied.

- **VOLUMES:** There should be as many volume serial numbers in the volser list as there are key ranges. When a volume serial number is duplicated in the volser list, more than one key range is allocated space on the volume. When more than one key range is contained on a volume, UNIQUE cannot be coded for the cluster's data component.

When there are more volumes in the volser list than there are key ranges, the excess volumes are used for overflow records from any key range without consideration for key-range boundaries.

When there are fewer volumes in the volser list than there are key ranges, the excess key ranges are allocated on the last volume specified—UNIQUE cannot also be specified.

- **UNIQUE:** When UNIQUE is specified, each key range resides on its own volume in its own VSAM data space. Other key ranges for the cluster cannot also reside on the volume.
- **ORDERED:** There is a one-for-one correspondence between the volumes in the volser list and the key ranges: The first volume on the volume list contains the first key range, the second volume contains the second key range, and so on. If a volume cannot be allocated in the order specified by the volser list, your cluster definition job terminates with an error message.
- **KEYS:** The length of the key values must not exceed the key length specified in the KEYS parameter. The KEY parameter must be specified in the same component as the KEYS parameter.

Keys can contain 1 to 64 characters; 1 to 128 hexadecimal characters if coded as X'lowkey' b X'highkey'. All EBCDIC characters are allowed. Keys consisting of characters must be enclosed in single quotation marks if they contain commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in single quotation marks.

lowkey

specifies the low key of the key range. If lowkey is shorter than the actual keys, it will be padded on the right with binary zeros.

highkey

specifies the high key of the key range. If highkey is shorter than the actual keys, it will be padded on the right with binary ones.

Abbreviation: KRNG

KEYS(length b offset|64 0)

specifies information about the key field of a key-sequenced data set's data records. This parameter applies only to key-sequenced clusters. The default specifies a key field of 64 bytes in length, beginning at the first byte (byte 0) of each data record.

length b offset

specifies the length of the key and its displacement (in bytes) from the beginning of the record. The sum of length + offset cannot exceed the length of the shortest record.

The length of the key can be from 1 through 255 bytes. length and offset can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

The key field of the cluster's index is called the prime key to distinguish it from other key fields, called alternate keys. See DEFINE ALTERNATEINDEX for more details on how to specify alternate keys for a cluster.

KEYS cannot be specified when you define an entry-sequenced cluster or a relative-record cluster. When the data record is allowed to span control intervals, the record's key field must be within the part of the record that is in the first control interval.

MASTERPW(password)

specifies a master password for the entry being defined. The master password allows all Access Method Services operations against the cluster entry and its data and index entries. The master password allows the user's program to access the cluster's contents without restriction. For more details on how passwords can be used, see "Data Security and Protection." The AUTHORIZATION, CODE, and ATTEMPTS parameters have no effect unless the entry has a master password associated with it. If MASTERPW is not specified, the highest-level password specified becomes the password for all higher levels.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation: MRPW

MODEL(entryname[/password] [b catname[/password]])

specifies that an existing entry is to be used as a model for the entry being defined. You can use an existing cluster's entry as a model for the attributes of the cluster being defined. For details about how a model is used, see "How to Use One Object as a Model for Another Object."

You may use some attributes of the model and override others by explicitly specifying them in the definition of the cluster or component. If you do not want to add or change any attributes, you need specify only the entry-type (cluster, data, or index) of the model to be used and the name of the entry to be defined.

When you use a cluster entry as a model for the cluster, the model-cluster's data and index entries are used as models for the to-be-defined cluster's data and index components, unless another entry is specified with the MODEL parameter as a subparameter of DATA or INDEX.

entryname

specifies the name of the cluster or component entry to be used as a model.

password

specifies a password. If the model entry is password protected and it is cataloged in a password-protected catalog, you must supply the read- (or higher-level) password of either the model entry or its catalog. If both passwords are supplied, the catalog's password is used.

If you are not specifying new protection attributes for the cluster (that is, the model's passwords and protection attributes are being copied), you must supply the master password of either the model entry or its catalog.

catname

names the model-entry's catalog. You identify the catalog that contains the model entry for either of these cases:

- You specify the catalog's password instead of the model-entry's password.
- The model-entry's catalog is not identified with a JOBCAT or STEPCAT DD statement, and is not the master catalog.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved. See "Order of Catalog Use: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

ORDERED|UNORDERED

specifies whether volumes are to be used in the order in which they were listed in the VOLUMES parameter. If KEYRANGES is also specified, all of the records within the range specified by the first low-key/high-key pair are placed on the first volume specified in VOLUMES; all of the records within the second range are placed on the second volume; etc. If it is impossible to allocate volumes in the given order and ORDERED is specified, the command is terminated.

Abbreviations: ORD and UNORD

OWNER(ownerid)

specifies the identification of the cluster's owner.

ownerid

can contain 1 to 8 alphameric or special characters.

If the ownerid contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*').)

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('X'CA'ROLX')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

Note to TSO Users: If the owner is not identified with the OWNER parameter, the TSO user's userid becomes the ownerid.

READPW(password)
specifies a read password for the entry being defined. The read password permits read operations against the cluster's data records.

password
is a 1 to 8 alphanumeric special character password.

If the password contains a lowercase alphabetic character it is folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('X'DORISX')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('X'CA'ROLX')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

RECORDSIZE(average b maximum|default)
specifies the average and maximum lengths, in bytes, of the records in the data component.

RECORDSIZE can be specified as a parameter of either CLUSTER or DATA.

average and maximum are integer values and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. The minimum record size that you can specify is 1 byte. For nonspanned records, the maximum recordsize + 7 cannot exceed the data-component's control interval size (that is, the maximum nonspanned record size, 32,761, + 7 equals the maximum data-component control interval size, 32,768). For spanned records, the maximum record size cannot exceed the number of control intervals per control area x (control interval size less ten), as calculated by VSAM. When you specify a record size that is larger than one control interval, you must also specify spanned records (SPANNED).

You can identify the records as fixed-length by specifying the same value for average and maximum. If NUMBERED is also specified, the records must be fixed-length (that is, average equals maximum).

When your records are fixed-length, you can use the following formula to find a control-interval size that contains a whole number (n) of records:

$$\text{CISZ} = (n \times \text{RECSZ}) + 10$$

or

$$n = (\text{CISZ} - 10) / \text{RECSZ}$$

where:

- n is the number of fixed-length records in a control interval, and is a positive integer.
- CISZ is the control interval size (see also the CONTROLINTERVALSIZE parameter).
- RECSZ is the average record size.

default

When SPANNED is specified, the default is RECORDSIZE(4086 32600). Otherwise, the default is RECORDSIZE(4089 4089).

Caution: When you specify RECORDS, you should ensure that:

$$\text{REC}(\text{sec}) \times \text{RECSZ}(\text{avg}) > \text{RECSZ}(\text{max})$$

where:

- REC(sec) is the secondary space allocation quantity, in records.
- RECSZ(avg) is the average record size (default = 4086 or 4089 bytes).
- RECSZ(max) is the maximum record size (default = 4089 or 32600 bytes).

When the SPANNED record size default prevails (32600 bytes), the secondary allocation quantity should be at least 8 records.

If you plan to copy your cluster using either the EXPORT or EXPORTRA command, you should be aware of the restriction in maximum record length imposed by these commands. Refer to the sections "Moving Entries" and "Copying a Catalog Entry From the Catalog Recovery Area" for details of this restriction.

Abbreviation: RECSZ

REPLICATE|NOREPLICATE

specifies whether each index record is to be written on a track as many times as it will fit. When you specify REPLICATE, rotational delay is reduced and performance is improved. However, the cluster's index usually requires more direct-access device space. This parameter applies only to key-sequenced clusters.

Refer to the description of the IMBED|NOIMBED parameter for a discussion of the relationship between IMBED|NOIMBED and REPLICATE|NOREPLICATE.

Abbreviations: REPL and NREPL

REUSE|NOREUSE

specifies whether the cluster can be opened again and again as a temporary, or reusable, cluster (that is, with its high-used RBA set to zero when you open it with an ACB which specifies the RESET attribute). REUSE allows you to create an entry-sequenced, key-sequenced, or relative-record workfile.

When you create a reusable cluster, you cannot build an alternate index to support it. Also, you cannot create a reusable cluster with key ranges (see the KEYRANGE parameter) or with its own data space (see the UNIQUE parameter). Reusable data sets may be multivolumed and are restricted to sixteen physical extents per volume.

Abbreviations: RUS and NRUS

SHAREOPTIONS(crossregion[b crosssystem][1 3])
specifies how a component or cluster can be shared among users. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for a description of data set sharing.

Unless you specify sharing options, the cluster cannot be shared.

crossregion

specifies the amount of sharing allowed among regions. Independent job steps in an operating system can access a VSAM data set concurrently. To share a data set, each user must specify DISP=SHR in the data set's DD statement. The values that can be specified are:

- 1 specifies that the data set can be shared by any number of users for read processing, or the data set can be accessed by only one user for read and write processing. With this option, VSAM ensures complete data integrity for the data set.
- 2 specifies that the data set can be accessed by any number of users for read processing and that it can also be accessed by one user for write processing. With this option, VSAM ensures write integrity by obtaining exclusive control for a control interval when it is to be updated. If a user desires read integrity, it is necessary to use the ENQ and DEQ macros appropriately to provide read integrity for the data the program obtains.
- 3 specifies that the data set can be fully shared by any number of users. With this option, each user is responsible for maintaining both read and write integrity for the data the program accesses. User programs that ignore the write integrity guidelines can cause VSAM program checks, lost or inaccessible records, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared by any number of users and buffers used for direct processing are refreshed for each request. This option requires your program to use the ENQ and DEQ macros to maintain data integrity while sharing the data set. Improper use of the ENQ macro can cause problems similar to those described under SHAREOPTION 3.

crosssystem

specifies the amount of sharing allowed among systems. Job steps of two or more OS/VS operating systems can gain access to the same VSAM data set regardless of the disposition specified in each step's DD statement for the data set. To get exclusive control of the data set's volume, a task in one system issues the RESERVE macro. The level of cross-system sharing allowed by VSAM applies only in a multiple operating system

environment. The values that can be specified are:

- 1 Reserved
- 2 Reserved
- 3 specifies that the data set can be fully shared. With this option, each user is responsible for maintaining both read and write integrity for the data the program accesses. User programs that ignore write-integrity guidelines can cause VSAM program checks, uncorrectable data set failures, and other unpredictable results. This option places heavy responsibility on each user sharing the data set.
- 4 specifies that the data set can be fully shared. Buffers used for direct processing are refreshed for each request. This option requires that you use the RESERVE and RELEASE macros to maintain data integrity while sharing the data set. Writing is limited to PUT-update and PUT-insert processing that does not change the high-used RBA if your program opens the data set with DISP=SHR. Data set integrity cannot be maintained unless all jobs accessing the data set in a cross-system environment specify DISP=SHR. Improper use of the RESERVE macro can cause problems similar to those described under SHAREOPTION 3.

The assistance VSAM provides each user of SHAREOPTION 4 to ensure data integrity in a shared environment is:

- Each PUT request results in the appropriate buffer(s) being written immediately to the VSAM cluster's direct-access device space (that is, the buffer in the user's address space that contains the new or updated data record, and the buffers that contain new or updated index records when the user's data is key-sequenced.)
- Each GET request results in all of the user's input buffers being refreshed. That is, each buffer's contents (each data and index buffer being used by the user's program) is retrieved from the VSAM cluster's direct access device.

To ensure integrity you should be sure that share options specified at the DATA and INDEX levels are the same.

Additional information about shared data can be found in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.

Abbreviation: SHR

SPANNED | NONSPANNED

specifies whether a data record is allowed to cross control-interval boundaries. When the maximum length of a data record (as specified with RECORDSIZE) is larger than a control interval, the record can be contained on more than one control interval. To ensure that you want this to happen, specify SPANNED. (This allows VSAM to select a control interval size that is optimum for the direct-access device.) Otherwise, VSAM would select a larger control interval size that accommodates your largest record.

When a data record which is larger than a control interval is put into a cluster that allows spanned records, the first part of the record completely fills a control interval. Subsequent control intervals are filled until the record is written into the cluster. Unused space in the record's last control interval is not available to contain other data records.

Abbreviations: SPND and NSPND

SPEED|RECOVERY

specifies whether storage allocated to the data component is to be preformatted before records are inserted. SPEED|RECOVERY applies only to initial loading.

When RECOVERY is specified, the data component's control areas are written with records that indicate end-of-file. When a data record is written (during the initial load) into a control interval, it is always followed by a record that identifies the just-written record as the last record in the cluster. If the initial load fails, you can resume loading data records after the last correctly-written data record (because an end-of-file indicator identifies it as the last record—that is, no more data records follow).

When SPEED is specified, the data component's space is not preformatted. Its space might contain data records from a previous use of the space, or it might contain binary zeros—its contents are unpredictable. If the initial load fails, you must load the data records again from the beginning (because VSAM is unable to determine where your last correctly-written record is—VSAM cannot find a valid end-of-file indicator when it searches your data records).

When you specify RECOVERY, your initial load takes longer because the control areas are written initially with end-of-file indicators, and again with your data records. When you specify SPEED, your initial load is quicker.

Abbreviation: RCVY

STAGE|BIND|CYLINDERFAULT

specifies how a cluster or component that is stored on a mass storage volume is to be staged.

STAGE

indicates that the cluster or component is to be staged from mass storage to a direct access storage staging drive when the cluster or component is opened. If the cluster or component cannot be staged at open time because of heavy staging activity of other objects, data is staged as a processing program needs it.

BIND

indicates that the cluster or component is not only to be staged, but also to be bound—that is, retained on the direct-access storage staging drive until it is closed. If the cluster or component cannot be staged at open time because of heavy staging activity of other objects, data is staged as a processing program needs it.

CYLINDERFAULT

indicates that the cluster or component is not to be staged when it is opened, but that data from it is to be staged as a processing program needs it.

Abbreviation: CYLF

When the cluster or component is not stored on a mass storage volume, the attribute is ineffective until the direct access storage volume the cluster or component is stored on is converted to a mass storage volume (by way of

the CONVERTV command, which is described in OS/VS Mass Storage System (MSS) Services: Reference Information).

When one of these parameters is specified for the data component and another parameter is specified for the index component, the components are staged separately as specified, except when the sequence set of the index component is imbedded in the data. In that case, the parameter specified for the index component applies to both components.

TO(date)|FOR(days)

specifies the retention period for the cluster being defined. If neither TO nor FOR is specified, the cluster can be deleted at any time.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the number (001 through 365) of the day, through which the cluster being defined is to be kept.

FOR(days)

specifies the number of days for which the cluster being defined is to be kept. The maximum number that can be specified is 9999. If the number specified is 0 through 1830, the cluster is retained for the number of days specified; if the number is between 1831 and 9999, the cluster is retained through the year 1999.

days

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

UNIQUE|SUBALLOCATION

specifies whether the cluster's components are allocated space of their own or whether a portion of previously defined VSAM data space is to be used for each component.

If UNIQUE is specified, the components of the cluster are allocated space separately. Their names appear in the VTOC of the volume(s) within a format-1 DSCB.

If SUBALLOCATION is specified, the name of the data space, not of the component, appears in the VTOC. If SUBALLOCATION is coded, a data space must exist on the volume on which the cluster or components are to reside. When UNIQUE is specified as a parameter of CLUSTER, a data space is built and allocated to each component of the cluster.

The space-allocation attribute interacts with other DEFINE CLUSTER parameters. You should take care to ensure that the attribute you specify for the cluster is consistent with other attributes:

- REUSE: You cannot specify REUSE when you specify UNIQUE for a cluster or its components.
- KEYRANGES: When UNIQUE is specified, a data space is built and allocated for each key range. Each key range is on a separate volume.
- VOLUMES: When UNIQUE is not specified, VSAM data space must exist on the volume that is to contain the cluster's component. When UNIQUE is specified, and more than one volume is specified, VSAM must already own all the volumes except the first. If there is no VSAM space on a volume, you must execute a DEFINE SPACE CANDIDATE before your DEFINE UNIQUE.

Abbreviations: UNQ and SUBAL

UPDATEPW(password)

specifies an update password for the entry being defined. The update password permits read and write operations against the cluster's data records.

If a read password is the only password specified for the object (that is, it is the highest-level password), it propagates upward and becomes the password for all higher levels. If you specify a higher-level password and do not specify an update password, the update password is null.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, UPDATEPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, UPDATEPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, UPDATEPW('X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

WRITECHECK|NOWRITECHECK

specifies whether the cluster or component is to be checked by a machine action called write check when a record is written into it. If WRITECHECK is specified, a record is written and then read, without data transfer, to test for the data check condition.

Abbreviations: WCK and NWCK

Components (Cluster)

Attributes can be specified separately for the cluster's data and index components.

DATA(options)

specifies attributes of the data component of the cluster. You can specify separate attributes for the data component of any type of cluster (that is, entry-sequenced, key-sequenced, and relative-record).

Except for name, space allocation, and protection attributes, attributes specified as parameters of DATA override the same attributes specified as parameters of CLUSTER. Attributes specified as parameters of DATA also override the same attributes of a model cluster's data entry. Parameters that can be specified for the cluster's data component are:

- **ATTEMPTS, AUTHORIZATION, and CODE:** Specifies protection attributes for the data component when it is opened as a separate VSAM object by the user's program.
- **BUFFERSPACE:** Specifies the amount of buffer space that the user's program is to provide from its address space when the program opens the data component as a separate

VSAM object.

- **CONTROLINTERVALSIZE:** Specifies the size of the data component's control interval. The sizes you can specify are:

$$CISZ = (n \times 512) \text{ or } (n \times 2048)$$
 where n is a positive integer from 1 to 16.
- **CYLINDERS, RECORDS, or TRACKS:** You can specify an amount of space to be allocated to the data component. When you specify space as a parameter of DATA, however, you cannot also specify space for the cluster as a whole. See the description of CYLINDERS | RECORDS | TRACKS as parameters of CLUSTER.
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the data component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that closes it.
- **ERASE or NOERASE:** Specifies whether the data records are to be erased when the cluster is deleted.
- **EXCEPTIONEXIT:** Identifies an exception exit routine for the cluster's data component.
- **FILE:** Names the DD statement that describes a direct access device and volume that is to contain the data component. When more than one volume is to contain the data component, all volumes for that component must be the same device type.
- **FREESPACE:** Specifies the amount of space that is to be left free during loading and after any split of control intervals and control areas. This parameter is valid only for a key-sequenced cluster.
- **KEYRANGES:** Specifies that portions of the data component of a key-sequenced cluster are to be put on different volumes, and specifies the boundaries of each key range.
- **KEYS:** Describes the key field of a key-sequenced cluster's data record.
- **MASTERPW, CONTROLPW, UPDATEPW, or READPW:** Specifies passwords for the data component when it is opened as a separate VSAM object by the user's program.
- **MODEL:** Identifies a data entry that is to be used as a model for the data component only.
- **NAME:** Names the data component. If you do not specify a name for the data component, VSAM generates one (see "Defining a Cluster" for the format of the generated name).
- **ORDERED or UNORDERED:** When ORDERED is specified, the volumes listed with the VOLUMES parameter are allocated in the order listed.
- **OWNER:** Identifies the data component's owner, when the owner is different from the owner of the cluster as a whole.
- **RECORDSIZE:** Specifies the average and maximum sizes of the cluster's data records.
- **REUSE or NOREUSE:** Specifies whether the data component is to be reusable.

- **SHAREOPTIONS:** Specifies share options for the data component when it is opened as a separate VSAM object by the user's program.
- **SPEED or RECOVERY:** Specifies whether the data control intervals are to be preformatted with the end-of-file indicator before data records are inserted.
- **SPANNED or NONSPANNED:** Specifies whether a data record is allowed to cross control interval boundaries.
- **STAGE, BIND, or CYLINDERFAULT:** Specifies whether the data component, stored on a mass storage volume, is to be staged.
- **UNIQUE or SUBALLOCATION:** Specifies whether the data component is to occupy a separate VSAM data space or share an existing data space with other VSAM objects.
- **VOLUMES:** Identifies the volume that is to contain the data component, and identifies other volumes that can be used as overflow volumes for the data component. All volumes must be of the same device type.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the data records are to be checked by the direct access device when they are written into the data component.

DATA and its subparameters can be specified for any type of cluster. All of the DATA subparameters are described in more detail as parameters of the cluster as a whole. Certain of the above subparameters apply only to the data component of a key-sequenced cluster.

INDEX(options)

specifies attributes of the index component of a key-sequenced cluster. Except for name, space allocation, and protection attributes, attributes specified as parameters of INDEX override the same attributes specified as parameters of CLUSTER. Attributes specified as parameters of INDEX also override the same attributes of a model cluster's index entry.

Parameters that can be specified for the cluster's index component are:

- **ATTEMPTS, AUTHORIZATION, and CODE:** Specifies protection attributes for the index when it is opened as a separate VSAM object by the user's program.
- **CONTROLINTERVALSIZE:** Specifies the size of the index's control interval. The sizes you can specify are: 512, 1024, 2048, or 4096 bytes.
- **CYLINDERS, RECORDS, or TRACKS:** You can specify an amount of space to be allocated to the index component. When you specify one of the space-quantity parameters as a parameter of INDEX, you must also specify a space-quantity parameter as a parameter of DATA. You cannot also specify a quantity of space for the cluster as a whole.
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the index component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the user's program that closes it.
- **EXCEPTIONEXIT:** Identifies an exception exit routine for the index component.
- **FILE:** The DD statement describes a direct access device that is to contain the index. When more than one volume is to contain the index component, all volumes for that component must be the same device type.

- **IMBED or NOIMBED:** Specifies whether the sequence set records are to be written adjacent to the data.
- **MASTERPW, CONTROLPW, UPDATEPW, or READPW:** Specifies passwords for the index when it is opened as a separate VSAM object by the user's program.
- **MODEL:** Identifies an index entry that is to be used as a model for the index component only.
- **NAME:** Names the index component. If you do not specify a name for the index component, VSAM generates one (see "Defining a Cluster" for the format of the generated name).
- **ORDERED or UNORDERED:** When **ORDERED** is specified, the first volume in the volser list (of the **VOLUMES** parameter) is to contain the index. Other volumes are to be available as overflow volumes. Otherwise, any volume specified in the volser list can be used for the index.
- **OWNER:** Identifies the index's owner, when the owner identification is different from the owner of the cluster as a whole.
- **REPLICATE or NOREPLICATE:** Specifies whether the index's records are to be repeated around the track.
- **REUSE or NOREUSE:** Specifies whether the index component is to be reusable.
- **SHAREOPTIONS:** Specifies share options for the index when it is opened as a separate VSAM object by the user's program.
- **STAGE, BIND, or CYLINDERFAULT:** Specifies how the index component, stored on a mass storage volume, is to be staged.
- **UNIQUE or SUBALLOCATION:** Specifies whether the index is to occupy a separate VSAM data space or share an existing data space with other VSAM objects.
- **VOLUMES:** Identifies the volume that is to contain the index, and identifies other volumes that can be used as overflow volumes for the index. All volumes must be of the same device type.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the index records are to be checked by the direct-access device when they are written into the index.

INDEX and its subparameters can be specified only for a key-sequenced cluster. Other restrictions are noted with each parameter's description. All of the **INDEX** subparameters are described in more detail as parameters of the cluster as a whole.

Catalog (Cluster)

The catalog parameter is used to supply the name and password, when required, of the catalog in which the cluster is to be defined.

CATALOG(catname[/password])

identifies the catalog in which the cluster is to be defined. See "Order of Catalog Use: DEFINE" for information about the order in which catalogs are used.

catname

specifies the name of the catalog in which the entry is to be defined.

password

specifies the catalog's password. If the catalog is password protected, you must supply the update- or higher-level password. If no password is specified, VSAM asks the operator for the correct password.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved.

Abbreviation: CAT

DEFINE CLUSTER EXAMPLES

Define a Key-Sequenced Cluster: Example 1

In this example, a key-sequenced cluster is defined. The DATA and INDEX parameters are specified and the cluster's data and index components are explicitly named. This example assumes that VSAM data space exists on volume VSER02. It also assumes that an alias name D40 has been defined for the catalog D27UCAT1. This naming convention causes D40.MYDATA to be cataloged in D27UCAT1.

```
//DEFCLU2 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
        DEFINE CLUSTER -
            (NAME(D40.MYDATA) -
             VOLUMES(VSER02) -
             RECORDS(1000 500)) -
        DATA -
            (NAME(KSDATA) -
             KEYS(15 0) -
             RECORDSIZE(250 250) -
             FREESPACE(20 10) -
             BUFFERSPACE(25000) ) -
        INDEX -
            (NAME(KSINDEX) -
             IMBED) -
        CATALOG (D27UCAT1/UPPWD27)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE CLUSTER command builds a cluster entry, a data entry, and an index entry to define the key-sequenced cluster D40.MYDATA. The parameters specified for the cluster as a whole are:

- NAME, which specifies that the cluster's name is D40.MYDATA.
- VOLUMES, which specifies that the cluster is to reside on volume VSER02.
- RECORDS, which specifies that the cluster's space allocation is 1,000 data records. When the cluster is extended, it is extended in increments of 500 records. After the space is allocated, VSAM calculates the amount required for the index and subtracts it from the total.

In addition to the parameters specified for the cluster as a whole, DATA and INDEX subparameters specify values and attributes that apply only to the cluster's data or index component. The parameters specified for the data component of D40.MYDATA are:

- NAME, which specifies that the data component's name is KSDATA.
- KEYS, which specifies that the length of the key field is 15 bytes and that the key field begins in the first byte (byte 0) of each data record.
- RECORDSIZE, which specifies fixed-length records of 250 bytes.
- BUFFERSPACE, which specifies that a minimum of 25000 bytes must be provided for I/O buffers. A large area for I/O buffers can help to improve access time with certain types

of processing. For example, with direct processing, if the high-level index can be kept in virtual storage, access time is reduced. With sequential processing, if enough I/O buffers are available, VSAM can perform a read-ahead thereby reducing system overhead and minimizing rotational delay.

- FREESPACE, which specifies that 20 percent of each control interval and 10 percent of each control area are to be left free when records are loaded into the cluster. After the cluster's records are loaded, the free space can be used to contain new records.

The parameters specified for the index component of D40.MYDATA are:

- NAME, which specifies that the index-component's name is KSINDEX.
- IMBED, which specifies that sequence-set index records are to be placed in the data component's control areas (the sequence-set records will be replicated automatically).

The parameter which provides the catalog's password is:

- CATALOG, which specifies the catalog name and its update password. This parameter would not be necessary if (1) the catalog was not password protected and (2) D40 has been defined as the catalog's alias.

Define a Key-Sequenced Cluster and an Entry-Sequenced Cluster: Example 2

In this example, two VSAM clusters are defined. The first DEFINE command defines a key-sequenced VSAM cluster, D40.EXAMPLE.KSDS1. The second DEFINE command defines an entry-sequenced VSAM cluster, D50.EXAMPLE.ESDS1. In both examples it is assumed that alias names, D40 and D50, have been defined for user catalogs D27UCAT1 and D27UCAT2, respectively, and that neither user catalog is password protected. This example assumes that VSAM data space that can contain the data sets already exists on volumes VSER02 and VSER03.

```
//DEFCLU2 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE CLUSTER -
              (NAME(D40.EXAMPLE.KSDS1) -
              MODEL(D40.MYDATA) -
              VOLUMES(VSER02) -
              NOIMBED )

          DEFINE CLUSTER -
              (NAME(D50.EXAMPLE.ESDS1) -
              RECORDS(100 500) -
              RECORDSIZE(250 250) -
              VOLUMES(VSER03) -
              NONINDEXED )
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The first DEFINE command builds a cluster entry, a data entry, and an index entry to define the key-sequenced cluster D40.EXAMPLE.KSDS1. Its parameters are:

- NAME, which specifies the name of the key-sequenced cluster, D40.EXAMPLE.KSDS1. The cluster will be defined in the user catalog for which D40 has been established as an alias.

- MODEL, which identifies D40.MYDATA as the cluster to use as a model for D40.EXAMPLE.KSDS1. The attributes and specifications of D40.MYDATA that are not otherwise specified with the DEFINE command's parameters are used to define the attributes and specifications of D40.EXAMPLE.KSDS1. D40.MYDATA is located in the user catalog for which D40 has been established as an alias.
- VOLUMES, which specifies that the cluster is to reside on volume VSER02.
- NOIMBED, which specifies that space is not to be allocated for sequence-set control intervals within the data-component's physical extents.

The second DEFINE command builds a cluster entry and a data entry to define an entry-sequenced cluster D50.EXAMPLE.ESDS1. Its parameters are:

- NAME, which specifies the name of the entry-sequenced cluster, D50.EXAMPLE.ESDS1. The cluster will be defined in the user catalog for which D50 has been established as an alias.
- RECORDS, which specifies that the cluster's space allocation is 100 records. When the cluster is extended, it is extended in increments of 500 records.
- RECORDSIZE, which specifies that the cluster's records are fixed-length (the average record size equals the maximum record size) and 250 bytes long.
- VOLUMES, which specifies that the cluster is to reside on volume VSER03.
- NONINDEXED, which specifies that the cluster is to be an entry-sequenced cluster.

Define a Key-Sequenced Cluster (In a Unique Data Space): Example 3

In this example, a key-sequenced cluster is defined. The cluster is unique; that is, it is the only cluster in a VSAM data space.

```
//DEFCLU3 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//VOL4   DD       VOL=SER=VSER04,UNIT=2314,DISP=OLD
//SYSIN  DD       *
          DEFINE CLUSTER -
            (NAME(ENTRY) -
            RECORDSIZE(80 80) -
            FILE(VOL4) -
            KEYS(10 10) -
            VOLUMES(VSER04) -
            UNIQUE -
            CYLINDERS(5 10) ) -
          CATALOG(MYCAT)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE CLUSTER command builds a cluster entry, a data entry, and an index entry to define the key-sequenced cluster ENTRY. The DEFINE CLUSTER command also allocates a data space and allocates it for the cluster's exclusive use. The command's parameters are:

- NAME, which specifies the cluster's name is ENTRY.
- RECORDSIZE, which specifies that the records are fixed-length, 80 byte records.
- FILE, which specifies the name of a DD statement that describes the cluster's volume and causes it to be mounted. The volume must be mounted so that VSAM can record the format-1 DSCBs of the key-sequenced cluster in the volume's VTOC. If FILE is not specified, the volume is dynamically allocated.
- KEYS, which specifies that the length of the key field is 10 bytes and that the key field begins in the 11th byte (byte 10) of each data record.
- VOLUMES and UNIQUE, which specify that ENTRY is to reside alone in a data space on volume VSER04. Access Method Services will create two data spaces for the cluster: one for the data component, and one for the index. Both data spaces are to be on volume VSER04. This example assumes that volume VSER04 has enough available space to contain the new data space. This example also assumes that either the volume's entry is in the MYCAT catalog or that volume VSER04 is not owned by a VSAM catalog at the beginning of the job.
- CYLINDERS, which specifies that five cylinders are allocated for the cluster's data space. When the cluster's data or index component is extended, the component is to be extended in increments of ten cylinders.
- CATALOG, which specifies the name of the catalog into which the cluster will be defined. In this example, MYCAT is not password protected and will be dynamically allocated.

Define a Relative-Record Cluster: Example 4

In this example, a relative-record cluster is defined. The cluster is suballocated (that is, it can reside in a VSAM data space with other VSAM objects). Volume VSER04 does not have to be mounted or allocated at this time.

```
//DEFCLU4 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE CLUSTER -
              (NAME(EXAMPLE.RRDS1) -
              RECORDSIZE(100 100) -
              VOLUMES(VSER04) -
              TRACKS(10 5) -
              NUMBERED) -
              CATALOG(AMASTCAT/MCATUPPW)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE CLUSTER command builds a cluster entry and a data entry to define the relative-record cluster EXAMPLE.RRDS1 in the master catalog. The DEFINE CLUSTER command also obtains space for the cluster from one of the VSAM data spaces on volume VSER04, and allocates ten tracks for the cluster's use. The command's parameters are:

- NAME, which specifies that the cluster's name is EXAMPLE.RRDS1.

- RECORDSIZE, which specifies that the records are fixed-length, 100 byte records. Average and maximum record length must be equal for a relative-record data set.
- VOLUMES, which specifies that the cluster is to reside on volume VSER04. This example assumes that the volume is already cataloged in the master catalog, AMASTCAT.
- TRACKS, which specifies that ten tracks are allocated for the cluster's space. When the cluster is extended, it is to be extended in increments of 5 tracks.
- NUMBERED, which specifies that the cluster's data organization is to be relative-record. This parameter overrides the INDEXED default.
- CATALOG, which supplies the master-catalog's update password, MCATUPPW.

Define a Reusable Entry-Sequenced Cluster: Example 5

In this example, a reusable entry-sequenced cluster is defined. The cluster can be used as a temporary data set. Each time the cluster is opened, its high-used RBA can be reset to zero. The cluster is suballocated (that is, it can reside in a VSAM data space with other VSAM objects).

```
//DEFCLUS JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE CLUSTER -
              (NAME(EXAMPLE.ESDS2) -
              RECORDSIZE(2500 3000) -
              SPANNED -
              VOLUMES(VSER03) -
              CYLINDERS(2 1) -
              NONINDEXED -
              REUSE -
              MASTERPW(ESD2MRPW) -
              UPDATEPW(ESD2UPPW) ) -
              CATALOG(D27UCAT2/UPDD27)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services jobs. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE CLUSTER command builds a cluster entry and a data entry to define the entry-sequenced cluster EXAMPLE.ESDS2. The DEFINE CLUSTER command also obtains space for the cluster from one of the Class 1 VSAM data spaces on volume VSER03, and assigns ten tracks for the cluster's use. VSER03 does not have to be mounted or allocated at this time. The command's parameters are:

- NAME, which specifies that the cluster's name is EXAMPLE.ESDS2.
- RECORDSIZE, which specifies that the records are variable-length, with an average size of 2500 bytes and a maximum size of 3000 bytes.
- SPANNED, which specifies that data records can cross control-interval boundaries.
- VOLUMES, which specifies that the cluster is to reside on volume VSER03. This example assumes that the volume is already cataloged in the user catalog, D27UCAT2.

- CYLINDERS, which specifies that two cylinders are allocated for the cluster's space. When the cluster is extended, it is to be extended in increments of one cylinder.
- NONINDEXED, which specifies that the cluster's data organization is to be entry-sequenced. This parameter overrides the INDEXED parameter.
- REUSE, which specifies that the cluster is to be reusable. Each time the cluster is opened, its high-used RBA can be reset to zero and it is effectively an empty cluster.
- MASTERPW and UPDATEPW, which specify the master password, ESD2MRPW, and the update password, ESD2UPPW, for the cluster.
- CATALOG, which specifies that the cluster is to be defined in a user catalog, D27UCAT2. The example also supplies the user-catalog's master password, UPDD27.

Define a Key-Sequenced Cluster: Example 6

In this example, a key-sequenced cluster is defined. In other examples, an alternate index is defined over the cluster, and a path is defined that relates the cluster to the alternate index. The cluster, its alternate index, and the path entry are all defined in the same catalog, AMASTCAT.

```
//DEFCLU6 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE CLUSTER -
            (NAME(EXAMPLE.KSDS2)) -
            DATA -
              (MASTERPW(DAT2MRPW) -
              UPDATEPW(DAT2UPPW) -
              READPW(DAT2RDPW) -
              RECORDS(500 100) -
              EXCEPTIONEXIT(DEXIT) -
              ERASE -
              FREESPACE(20 10) -
              KEYS(6 4) -
              RECORDSIZE(80 100) -
              VOLUMES(VSER04) ) -
            INDEX -
              (MASTERPW(IND2MRPW) -
              UPDATEPW(IND2UPPW) -
              READPW(IND2RDPW) -
              RECORDS(300 300) -
              IMBED -
              VOLUMES(VSER04) ) -
            CATALOG(AMASTCAT/MCATUPPW)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE CLUSTER command builds a cluster entry, a data entry, and an index entry to define the key-sequenced cluster EXAMPLE.KSDS2. The DEFINE CLUSTER command also obtains space for the cluster from one of the VSAM data spaces on volume VSER04, and allocates space separately for the cluster's data and index components. Since the cluster is being defined into a recoverable catalog, the catalog recovery area will be dynamically allocated on VSER04.

The command's parameter that applies to the cluster is:

- NAME, which specifies that the cluster's name is EXAMPLE.KSDS2.

The command's parameters that apply only to the cluster's data-component are enclosed in the parentheses following the DATA keyword:

- MASTERPW, UPDATEPW, and READPW, which specify the data component's master password, DAT2MRPW, update password, DAT2UPPW, and read password, DAT2RDPW.
- RECORDS, which specifies that an amount of tracks equal to at least 500 records is to be allocated for the data component's space. When the data component is extended, it is to be extended in increments of tracks equal to 100 records.
- EXCEPTIONEXIT, which specifies the name of the exception exit routine, DATEXIT, that is to be processed if an I/O error occurs while a data record is being processed.
- ERASE, which specifies that the cluster's data is to be erased (overwritten with binary zeros) when the cluster is deleted.
- FREESPACE, which specifies the amounts of free space to be left in the data component's control intervals (20%) and the control areas (10% of the control intervals in the control area) when data records are loaded into the cluster.
- KEYS, which specifies the location and length of the key field in each data record. The key field is 6 bytes long and begins in the fifth byte (byte 4) of each data record.
- RECORDSIZE, which specifies that the cluster's records are variable-length, with an average size of 80 bytes and a maximum size of 100 bytes.
- VOLUMES, which specifies that the cluster is to reside on volume VSER04. This example assumes that the volume is already cataloged in the master catalog, AMASTCAT.

The command's parameters that apply only to the cluster's index component are enclosed in the parentheses following the INDEX keyword:

- MASTERPW, UPDATEPW, and READPW, which specify the index-component's master password, IND2MRPW, update password, IND2UPPW, and read password, IND2RDPW.
- RECORDS, which specifies that an amount of tracks equal to at least 300 records is to be allocated for the index-component's space. When the index component is extended, it is to be extended in increments of tracks equal to 300 records.
- IMBED, which specifies that the index's sequence set records are to be placed in the data-component's control areas (the sequence set records will be replicated automatically).
- VOLUMES, which specifies that the index component is to reside on volume VSER04.

The CATALOG parameter supplies the catalog's update password.

Define an Entry-Sequenced Cluster (The Cluster has a Generic Name): Example 7

In this example, two entry-sequenced clusters are defined. Each has the generic name "GENERIC.*.BAKER," where the asterisk (*) is replaced with a simple name that uniquely identifies each cluster.

```

//DEFCLU4 JOB      ...
//STEP 1 EXEC     PGM=IDCAMS
//STEP CAT DD     DSN= D27UCAT1, DISP=OLD
//SYS PRINT DD    SYSOUT=A
//SYS IN DD       *
      DEFINE CLUSTER -
          (NAME(GENERIC.A.BAKER) -
           VOLUMES(VSER02) -
           RECORDS(100 100) -
           RECORDSIZE(80 80) -
           NONINDEXED ) -
          CATALOG(D27UCAT1/MRPWD27)

      DEFINE CLUSTER -
          (NAME(GENERIC.B.BAKER) -
           MODEL(GENERIC.A.BAKER D27UCAT1)) -
          CATALOG(D27UCAT1/MRPWD27)
/*

```

The job control language statements are:

- STEPCAT DD, which makes a catalog available for this job step: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The first DEFINE CLUSTER command defines an entry-sequenced cluster, GENERIC.A.BAKER. Its parameters are:

- NAME, which specifies the name of the entry-sequenced cluster, GENERIC.A.BAKER.
- VOLUMES, which specifies that the cluster is to reside on volume VSER02.
- RECORDS, which specifies that the cluster's space allocation is 100 records. When the cluster is extended, it is extended in increments of 100 records.
- RECORDSIZE, which specifies that the cluster's records are fixed-length (the average record size equals the maximum record size) and 80 bytes long.
- NONINDEXED, which specifies that the cluster is entry-sequenced.
- CATALOG, which specifies that the cluster is to be defined in the D27UCAT1 catalog. The master password of D27UCAT1 is MRPWD27.

The second DEFINE CLUSTER command uses the attributes and specifications of the previously defined cluster, GENERIC.A.BAKER, as a model for the to-be-defined cluster, GENERIC.B.BAKER. Its parameters are:

- NAME, which specifies the name of the entry-sequenced cluster, GENERIC.B.BAKER.
- MODEL, which identifies GENERIC.A.BAKER, cataloged in user catalog D27UCAT1, as the cluster to use as a model for GENERIC.B.BAKER. The attributes and specifications of GENERIC.A.BAKER that are not otherwise specified with the DEFINE command's parameters are used to define the attributes and specifications of GENERIC.B.BAKER.
- CATALOG, which specifies that the cluster is to be defined in the D27UCAT1 catalog. The master password of D27UCAT1 is MRPWD27.

DEFINE GENERATIONDATAGROUP

The format of the DEFINE command when it is used to define a generation data group is:

DEFINE	GENERATIONDATAGROUP (NAME(entryname) LIMIT(limit) [EMPTY <u>NOEMPTY</u>] [OWNER(ownerid)] [SCRATCH <u>NOSCRATCH</u>] [TO(date) FOR(days)]) [CATALOG(catname[/password])]
---------------	--

DEFINE may be abbreviated: DEF.

DEFINE GENERATIONDATAGROUP PARAMETERS

Required Parameters

GENERATIONDATAGROUP
 specifies that a generation-data-group entry is to be defined.

Abbreviation: GDG

NAME(entryname)
 specifies the name of the generation data group that is being defined. The name may consist of 1 through 35 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and the 12-0 overpunch). If the name is greater than eight characters, it is specified as a qualified name. A qualified name is segmented by periods; one to eight characters can be specified between periods. The first character of any name or name segment must be either an alphabetic or national character.

LIMIT(limit)
 specifies the maximum number, from 1 to 255, of generation data sets that can be associated with the generation data group to be defined. number can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviation: LIM

Optional Parameters

EMPTY|NOEMPTY
 specifies what action is to be taken when the maximum number of generation data sets has been reached for the generation data group and another generation-set data set is to be cataloged. The disposition of the data-set's DSCB in the volume's VTOC is determined with the SCRATCH|NOSCRATCH parameter.

EMPTY
 specifies that all the generation data sets are to be uncataloged when the limit is reached (that is, each data set's nonVSAM entry is automatically deleted from the catalog).

Abbreviation: EMP

NOEMPTY
 specifies that only the oldest generation data set is to be uncataloged when the limit is reached.

Abbreviation: NEMP

OWNER(ownerid)

specifies the identification of the owner of the generation data group being defined.

ownerid

The ownerid may contain 1 to 8 EBCDIC characters.

The ownerid must be enclosed in single quotation marks if it contains a comma, semicolon, blank, parenthesis, or slash. If a single quotation mark appears within ownerid, it must be coded as two single quotation marks when the ownerid is enclosed in single quotation marks.

Ownerid may be coded in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. If ownerid is specified in hexadecimal, it must be preceded by X and be enclosed in single quotation marks.

If Access Method Services is being used interactively with TSO to define a generation data group, and OWNER is not specified, the TSO userid is the default ownerid.

SCRATCH|NOSCRATCH

specifies whether a generation data set's DSCB is to be deleted from the volume's VTOC when the data set is uncataloged (that is, when its entry is deleted from the catalog automatically, as a result of EMPTY|NOEMPTY, or specifically, as a result of a user-issued DELETE request). The user can override the SCRATCH|NOSCRATCH attribute when issuing the DELETE command.

SCRATCH

specifies that the generation data-set's DSCB is to be deleted from the volume's VTOC when it is uncataloged. Direct-access device space management (DADSM) removes the data set's DSCB from the VTOC, erases the data-set's space on the volume, and makes the space available to other system users. The generation data set ceases to exist.

Note: On MSS volumes, the data set may not be scratched when it is automatically uncataloged, even when the SCRATCH option is specified.

Abbreviation: SCR

NOSCRATCH

specifies that the generation data set's DSCB is not to be removed from the volume's VTOC when it is uncataloged. The data set's DSCB in the volume's VTOC is intact and is used to locate the data set. The data set can no longer be accessed through the catalog. However, your program can process the data set as it processes any other nonVSAM data sets—that is, by using a JCL DD statement to describe the data set and allocate it.

Abbreviation: NSCR

TO(date)|FOR(days)

specifies the retention period for the generation data group being defined.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the number (001 through 365) of the day, through which the generation data group being defined is to be kept.

FOR(days)

specifies the number of days for which the generation data group being defined is to be kept. This value can be expressed in decimal, hexadecimal, or binary. If it is expressed in hexadecimal or binary, it must be preceded by X or B and be enclosed in single quotation marks. The maximum number that can be specified is 9999. If the number specified is 0 through 1830, the generation data group is retained for the number of days specified; if the number is between 1831 and 9999, the generation data group is retained through the year 1999. If neither TO nor FOR is specified, the generation data group can be deleted at any time.

Catalog (Generation Data Group)**CATALOG(catname[b password])**

identifies the catalog in which the generation data group is to be defined. If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved. See "Order of Catalog Search: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

catname

specifies the name of the catalog.

password

specifies the catalog's password. If the catalog is password protected, you must supply the update- or higher-level password.

Abbreviation: CAT

DEFINE GENERATIONDATAGROUP EXAMPLE

Define a Generation Data Group and a Generation Data Set Within It: Example 1

In this example, a generation data group is defined in the master catalog. Next, a generation data set is defined within the generation data group by using JCL statements.

```
//DEFGDG JOB ...
//STEP1 EXEC PGM=IDCAMS
//GDGMOD DD DSN=GDG01,DISP=(,KEEP),
// SPACE=(TRK,(0)),UNIT=2314,VOL=SER=VSER03,
// DCB=(RECFM=FB,BLKSIZE=2000,LRECL=100)
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
        DEFINE GENERATIONDATAGROUP -
        (NAME(GDG01) -
        EMPTY -
        NOSCRATCH -
        LIMIT(255) )
/*
//STEP2 EXEC PGM=IEFBR14
//GDGDD1 DD DSN=GDG01(+1),DISP=(NEW,CATLG),
// SPACE=(TRK,(10,5)),VOL=SER=VSER03,
// UNIT=2314
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
/*
```

The job control language statements are:

- GDGMOD DD, which describes the generation data group. When the scheduler processes the DD statement, no space is allocated to GDG01.

The model DSCB must exist on the generation data group's catalog volume.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

Because no catalog was specified in a CATALOG parameter or with a JOBCAT or STEPCAT DD statement, VSAM assumes all entries built for this command sequence are to be cataloged in the master catalog.

The DEFINE GENERATIONDATAGROUP command defines a generation data group base catalog entry, GDG01. Its parameters are:

- NAME, which specifies the name of the generation data group, GDG01. Each generation data set in the group will have the name GDG01.GxxxxVyy, where "xxxx" is the generation number and "yy" is the version number.
- EMPTY, which specifies that all data sets in the group are to be uncataloged by VSAM when the group contains the maximum number of data sets (as specified by the LIMIT parameter and one more generation data set is added to the group).
- NOSCRATCH, which specifies that, when a data set is uncataloged, its DSCB is not to be removed from its volume's VTOC. Therefore, even if a data set is uncataloged, its records can be accessed when it is allocated to a job step with the appropriate JCL DD statement.
- LIMIT, which specifies that the maximum number of generation data sets in the group is 255. The LIMIT parameter is required.

The second step, STEP2, is used to allocate space and catalog a generation data set in the newly defined generation data group. Its job control statements are:

- GDGDD1 DD, which specifies a generation data set in the generation data group.
- SYSPRINT DD, which is required in all job steps. The SYSPRINT DD statement identifies the output device to which messages to the programmer are to be sent.

DEFINE NONVSAM

The format of the DEFINE command when it is used to define a nonVSAM data set is:

DEFINE	NONVSAM (NAME(entryname) DEVICETYPES(devtype[b devtype...]) VOLUMES(volser[b volser...]) [FILESEQUENCENUMBERS(number[b number...])] [OWNER(ownerid)] [TO(date) FOR(days)]] [CATALOG(catname[/password])])
---------------	---

DEFINE may be abbreviated: DEF.

DEFINE NONVSAM PARAMETERS

Required Parameters

NONVSAM

specifies that a nonVSAM data set is to be defined.

Abbreviation: NVSAM

NAME(entryname)

specifies the name of the nonVSAM data set being defined. The entryname is the name that appears in the catalog; it is the name used in all future references to the data set. The entryname must be unique within the catalog in which it is defined.

Relative generation numbers (that is, GDGname(+1)) cannot be used with the entryname when you use the DEFINE NONVSAM command to catalog a generation data set and attach it to a generation data group. You identify a generation data set with its generation-data-group name (GDGname) followed by the data set's generation and version numbers: NAME(GDGname.GxxxxVyy).

The name may consist of 1 through 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and the 12-0 overpunch, X'C0'). Names containing more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of any name or name segment must be either an alphabetic or national character.

DEVICETYPES(devtypes[b devtype...])

is a required parameter that specifies the device types of the volumes containing the nonVSAM data set being defined. If the nonVSAM data set resides on different device types, the device types must be specified in the same order as the volume serial numbers listed in the VOLUMES parameter.

You can specify a device type for any device that is supported by your VS2 system. You can specify a generic device type (that is, 3330, 2314, etc.) or a specific unit address (that is, 121, 247, etc.). You can specify an esoteric device type (for example, SYSDA or TAPE) only if it has been established in the system device name table during system generation.

Warning: When an esoteric device type is specified, the device type put in the catalog record is the unique device type generated for that unit name in the system device name

table during SYSGEN. This unique device type is a function of the order of that unit name in the Stage I SYSGEN deck.

If the catalog is shared between systems, and if the unit names were not ordered the same in all Stage I SYSGEN decks, this device type may be different. An I/O GEN may also change this device type if the order of unit names in the Stage I SYSGEN deck is different. Allocation errors will result if a data set cataloged with an esoteric device type is referred to on a system in which the unit names were SYSGENed differently. See the discussion of the SCRATCH|NOSCRATCH option of the DELETE command for a restriction on deleting data sets that have been defined with an esoteric device type.

If the Data Facility/Device Support program product (5740-AM7) is installed on your system, and you expect to change the device type, you can code DEVICETYPES(0000) and this field will be resolved at SUPERLOCATE, LOCATE, and DELETE time to the device type. This will allow you to use the nonVSAM data sets without having to recatalog them to point to the new volume. When you code DEVICETYPES(0000) you must also code VOLUMES('*****'), or an error will result.

Abbreviation: DEVT

VOLUMES(volser[b volser...])

specifies the volumes to contain the nonVSAM data set. In a system with the Mass Storage System, a nonVSAM data set can be defined on a mass storage volume. You can specify more than one volume serial number if the data set resides on many volumes. If the data set resides on magnetic tape and more than one file belongs to the data set on a single tape volume, you must repeat the volume's serial number in order to maintain a one-to-one correspondence between the volume serial numbers and the file sequence numbers. For example, if your nonVSAM data set is contained in the first three files of magnetic tape volume TAPE10, you specify:

```
DEVICETYPES(2400 2400 2400) -
VOLUMES(TAPE10 TAPE10 TAPE10) -
FILESEQUENCENUMBERS(1 2 3) -
```

volser

can contain 1 to 6 alphameric or special characters.

If the volser contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character. When the volser contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the volser in single-quotation marks (for example, VOLUMES('*DORIS')).

When the volser contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, VOLUMES('*CA'ROL')). You can code volser in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, VOLUMES(X'E2E4C5') is the same as VOLUMES(SUE).

If the Data Facility/Device Support program product (5740-AM7) is installed on your system, and you expect to change the system residence volume serial number you can code VOLUMES('*****') and this field will be resolved at SUPERLOCATE, LOCATE, and DELETE time to the system

residence volume serial number. This will allow you to use the nonVSAM data sets without having to recatalog them to point to the new volume.

Abbreviation: VOL

Optional Parameters

FILESEQUENCENUMBERS(number[b number...])
specifies the file sequence number of the nonVSAM data set being defined. This number indicates the position of the file being defined with respect to other files of the tape. If the data set spans volumes, the file sequence number on each volume must be specified. The numbers must be specified in the same order as the volumes in the VOLUMES parameter.

For example, if your nonVSAM data set is contained in the first three files of magnetic tape volume TAPE10, you specify:

```
DEVICETYPES(2400 2400 2400) -  
VOLUMES(TAPE10 TAPE10 TAPE10) -  
FILESEQUENCENUMBERS(1 2 3) -
```

number
is a file sequence number. number can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviation: FSEQN

OWNER(ownerid)
specifies the identification of the nonVSAM data set.

ownerid
can contain 1 to 8 EBCDIC characters. When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('*CA'ROL*')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

If Access Method Services is being used interactively with TSO to define a nonVSAM data set, and OWNER is not specified, the TSO userid is the default ownerid.

TO(date)|FOR(days)
specifies the retention period for the nonVSAM data set being defined. The nonVSAM data set is not automatically deleted when the expiration date is reached. When you do not specify a retention period, the nonVSAM data set can be deleted at any time.

The maximum number that can be specified for days is 9999. If the number specified is 0 through 1830, the retention period is the number of days specified. If the number specified is between 1831 and 9999, the retention period is through the year 1999.

FOR(days)

specifies the number of days for which the nonVSAM data set is to be kept before it is allowed to be deleted. days can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the Julian date (001, for January 1, through 365, for December 31), through which the nonVSAM data set is to be kept before it is allowed to be deleted.

Catalog (NonVSAM)**CATALOG(catname[/password])**

identifies the catalog in which the nonVSAM data set is to be defined.

See "Order of Catalog Search: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

catname

specifies the name of the catalog in which the entry is to be defined.

password

specifies the catalog's password. If the catalog is password protected, you must supply the update- or higher-level password.

Abbreviation: CAT

DEFINE NONVSAM EXAMPLE

Define a NonVSAM Data Set: Example 1

In this example, two existing nonVSAM data sets are defined in a VSAM catalog, D27UCAT1. The DEFINE NONVSAM command cannot be used to create a nonVSAM data set because the command does not allocate space.

```
//DEFNVS JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE NONVSAM -
              (NAME(EXAMPLE.NONVSAM) -
               DEVICETYPES(2314) -
               VOLUMES(VSER02) ) -
              CATALOG(D27UCAT1/MRPWD27)

          DEFINE NONVSAM -
              (NAME(EXAMPLE.NONVSAM2) -
               DEVICETYPES(2314) -
               VOLUMES(VSER02) ) -
              CATALOG(D27UCAT1/MRPWD27)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

Both DEFINE NONVSAM commands define a nonVSAM data set in catalog D27UCAT1. The DEFINE NONVSAM commands' parameters are:

- NAME, which specifies the name of the nonVSAM data sets, EXAMPLE.NONVSAM and EXAMPLE.NONVSAM2.
- DEVICETYPES, which specifies the type of device that contains the nonVSAM data sets, 2314 Direct Access Storage Device.
- VOLUMES, which specify the volume that contains the nonVSAM data sets, VSER02.
- CATALOG, which identifies the catalog that is to contain the nonVSAM entries, D27UCAT1, and its update- or higher-level password, MRPWD27.

DEFINE PAGESPACE

The format of the DEFINE command when it is used to define a page space is:

DEFINE	PAGESPACE (NAME(entryname) {CYLINDERS(primary) RECORDS(primary) TRACKS(primary)} VOLUME(voiser) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLPW(password)] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [SWAP NOSWAP] [TO(date) FOR(days)] [UNIQUE SUBALLOCATION] [UPDATEPW(password)] [CATALOG(catname[/password])]
---------------	---

DEFINE PAGESPACE Parameters: Summary

The parameters of the DEFINE PAGESPACE can be divided into required and optional parameters.

The required parameters are:

- Entry type, which specifies that a page space and its catalog entry are to be created (PAGESPACE).
- Name, which names the page space (NAME).
- Allocation, which specifies the amount of space to be allocated to the page space and identifies its volume (VOLUMES and one of: CYLINDERS, RECORDS, or TRACKS). You need not explicitly specify the allocation parameters if you specify the MODEL parameter.

The optional parameters are:

- Those that apply only to the page-space data sets (SWAP and NOSWAP).
- Allocation, which specifies
 - Whether the page space's data space is to be shared with other VSAM objects (SUBALLOCATION and UNIQUE).
- Protection and integrity, which
 - Associate passwords with the page space (MASTERPW, CONTROLPW, UPDATEPW, READPW).
 - Provide a mechanism by which the console operator can be prompted to supply a password without disclosing the name of the entry (CODE).
 - Identify a user-written routine for additional authorization verification (AUTHORIZATION).
 - Identify the owner of the page space (OWNER).

- Specify a retention period for the page space (TO and FOR).
- Model, which identifies an existing page-space entry from which attributes are to be copied (MODEL).
- Catalog, which provides the name and password of the catalog in which the page space is to be defined (CATALOG).

DEFINE can be abbreviated: DEF.

DEFINE PAGESPACE PARAMETERS

Required Parameters

PAGESPACE

specifies that a page space is to be defined.

Abbreviation: PGSPC

NAME(entryname)

specifies the name of the entry being defined.

The name may contain from 1 through 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and the 12-0 overpunch). Names containing more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of any name or name segment must be either an alphabetic or national character. If you specify the name in the format entryname(modifier), only that portion of the name preceding the left parenthesis will be used. The modifier enclosed in parentheses will be ignored.

CYLINDERS(primary)|

RECORDS(primary)|

TRACKS(primary)

specifies the amount of space that is to be allocated by tracks, cylinders, or number of records. If RECORDS or TRACKS is specified, the quantity specified is rounded up to the nearest cylinder and the space is allocated in cylinders.

If you do not specify the MODEL parameter, you must specify one, and only one, of the following parameters: CYLINDERS, RECORDS, or TRACKS.

To determine the exact amount of space allocated, list the page space's catalog entry, using the LISTCAT command.

When you specify UNIQUE, and the page space's data space is the first data space on the volume that belongs to a recoverable catalog, an additional amount of (the equivalent, in tracks, of) one cylinder is allocated for the recovery area data space.

primary

specifies the amount of space that is to be allocated to the page space. Once the primary extent is full, the page space is full. The page space cannot extend onto secondary extents.

primary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

The Auxiliary Storage Manager (ASM) uses a halfword field to contain a relative paging slot (record) number within a given page space (data set). As such, the theoretical maximum number of paging slots for each page space is 65,535. This limitation currently affects only page space allocations on a 3350 device.

The maximum usable page space on a 3350 is 504 cylinders, which is 15,120 tracks, or 65,520 slots.

Abbreviations: CYL, REC, and TRK

VOLUMES(volser)

specifies the volume that contains the page space.

In a system with the Mass Storage System, a page space cannot reside on a mass storage volume. Also, a page space with the SWAP attribute cannot be defined on a 2314 direct access drive.

volser

can contain 1 to 6 alphameric or special characters.

If the volser contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the volser contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the volser in single-quotation marks (for example, VOLUMES ('*DORIS*')).

When the volser contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, VOLUMES ('*CA''ROL'.))

You can code volser in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, VOLUMES (X'E2E4C5') is the same as VOLUMES (SUE).)

You must specify the VOLUMES parameter unless you specify the MODEL parameter. The VOLUMES parameter interacts with other DEFINE PAGESPACE parameters. You should take care to ensure that the volume(s) you specify for the page space can satisfy the page space's other attributes:

- **SUBALLOCATION:** If suballocation is specified, the volume must contain a VSAM data space.
- **CYLINDERS, RECORDS, TRACKS:** The volume must contain enough unallocated space to satisfy the page space's space requirement.
- **FILE:** The volume information supplied with the DD statement pointed to by FILE must be consistent with the information specified for the page space.
- **CATALOG:** If the page space is suballocated, the data space on the volume must have been defined in the same catalog that the page space is being defined in, and must be owned by the catalog.

Abbreviation: VOL

Optional Parameters

ATTEMPTS(number|2)

specifies the maximum number of times the operator or TSO terminal user can try to enter a correct password in response to a prompting message. The number can be any number, 0 through 7. When you define a page space, you should specify ATTEMPTS(0), so that the operator is not prompted and is not allowed to enter a password from the console.

number

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. This parameter can be coded, but only has effect when the entry's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

Note to TSO users: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password for the page space. Checking the logon password counts as one attempt to obtain a password. If ATTEMPTS is not specified, the user has one attempt to supply the page-space's password because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])

specifies that a user-security-verification routine (USVR) is available for additional security verification. When a protected page space is accessed and the user supplies a correct password other than the page-space's master password, the USVR receives control.

See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for details on the USVR.

If a USVR is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

This parameter can be coded, but only has effect when the entry's master password is not null.

entrypoint

specifies the name of the user's-security-verification routine.

The name can contain 1 to 8 alphanumeric characters (A through Z, and 0 through 9), national characters (@, #, and \$), or special characters (the hyphen and 12-0 overpunch). The first character must be an alphabetic or national character.

string

specifies information to be passed to the user-security-verification routine when it receives control to verify authorization, and can contain 1 to 255 EBCDIC characters. When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single-quotation marks (for example, AUTHORIZATION(entrypoint, '*DORIS*')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, '*CA''ROL*')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint,SUE). The string can

contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

CODE(code)

specifies a code name for the page space. If an attempt is made to access a password-protected entry without a password, the code name is used in a prompting message; the code enables the operator or TSO terminal user to be prompted for the password without disclosing the name of the entry. If code is not specified and an attempt is made to access a page space entry that is password protected without supplying a password, the operator or TSO terminal user is prompted with the name of the entry.

This parameter can be coded, but only has effect when the cluster's or component's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

code

can contain 1 to 8 alphanumeric or special characters. When the code contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quotation mark ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the code in single-quotation marks (for example, CODE('*DORIS*')).

If the code contains a lowercase alphabetic character it is folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE('*CA'ROLX*')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, CODE(X'E2E4C5') is the same as CODE(SUE).

CONTROLPW(password)

specifies a control-level password for the page space. Because the page space is a system data set, it cannot be opened or used by a user's program.

If a read or update password is the only password specified for the page space, it (the highest-level password) propagates upward and becomes the password for all higher unspecified levels.

The passwords are cataloged in both the page-space entry and its data-component's entry. You should specify control-, update-, and read-level passwords to prevent the page space from being accessed if its volume is moved to a VS1 system. A VS2 system automatically prevents users from accessing page spaces and their data components.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

FILE(dname)

specifies the name of the DD statement that identifies the device and volume to be allocated to the page space. If the page space is unique, and if the FILE parameter is not specified and the volume is physically mounted, the volume identified with the VOLUME parameter is dynamically allocated. The volume must be mounted as permanently resident or reserved.

MASTERPW(password)

specifies a master password for the page space. The master password allows all Access Method Services operations against the page space entry. The AUTHORIZATION, CODE, and ATTEMPTS parameters have no effect unless the entry has a master password associated with it. If MASTERPW is not specified, the highest level password specified becomes the password for all higher levels. Because the page space is a system data set, it cannot be opened or used by a user's program.

The passwords are cataloged in both the page-space entry and its data-component's entry. You should specify control-, update-, and read-level passwords to prevent the page space from being accessed if its volume is moved to a VS1 system. A VS2 system automatically prevents users from accessing page spaces and their data components.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation: MRPW

MODEL(entryname[/password]
[b catname[/password]])

specifies that an existing page space entry is to be used as a model for the entry being defined. It is possible to use an already defined page space as a model for another page space. When one entry is used as a model for another, its attributes are copied as the new entry is defined.

For more details on how a model is used, see "How to Use One Object as a Model For Another Object."

You may use some attributes of the model and override others by explicitly specifying them in the definition of the page space. If you do not want to add or change any attributes, you need specify only the entry-type (page space) of the model to be used and the name of the entry to be defined.

entryname

specifies the name of the page space entry to be used as a model.

password

specifies a password. If the entry to be used as a model is password protected and is cataloged in a password-protected catalog, a password is required. If the protection attributes are to be copied, substitute the master password of either the entry being used as a model (following entryname) or the catalog in which the entry being used as a model is defined (following catname). If you specify both passwords, the catalog's password is used. If protection attributes are not to be copied, any password can be used.

catname

specifies the name of the catalog in which the entry to be used as a model is defined. You identify the catalog that contains the model entry for either of these cases:

- You specify the catalog's password instead of the model entry's password.
- The model entry's catalog is not identified with a JOBCAT or STEPCAT DD statement, and is not the master catalog.

OWNER(ownerid)

specifies the identification of the owner of the page space.

ownerid

can contain 1 to 8 EBCDIC characters.

When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('*CA'ROL*')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

Note to TSO users: If the owner is not identified with the OWNER parameter, the TSO user's userid becomes the ownerid.

READPW(password)

specifies a read level password for the page space. Because the page space is a system data set, it cannot be opened or used by a user's program.

The passwords are cataloged in both the page space entry and its data component's entry. You should specify control-, update-, and read-level passwords to prevent the page space from being accessed if its volume is moved to a VS1 system. A VS2 system automatically prevents users from accessing page spaces and their data components.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('*CA'ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

SWAP|NOSWAP

specifies for which group of pages the page space will be used. Auxiliary Storage Management separates private-area address space pages into two distinct groups: Local System Queue Area (LSQA) pages and pageable private-area pages. SWAP specifies that the page space is a high-speed data set used during a swap operation to store and retrieve the set of LSQA pages owned by an address space. NOSWAP indicates that the page space is a conventional page space used to record pageable private area pages.

A page space with the SWAP attribute cannot be defined on a 2314 direct access device.

Abbreviation: NSWAP

TO(date)|FOR(days)

specifies the retention period for the page space.

If neither TO nor FOR is specified, the page space can be deleted at any time.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the number (001 through 365) of the day, through which the cluster being defined is to be kept.

FOR(days)

specifies the number of days for which the page space is to be kept. The maximum number that can be specified is 9999. If the number specified is 0 through 1830, the page space is retained for the number of days specified; if the number is between 1831 and 9999, the page space is retained through the year 1999.

days

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

UNIQUE|SUBALLOCATION

specifies whether the page space is allocated an amount of space from the volume's available space (UNIQUE) or from a previously defined VSAM data space (SUBALLOCATION). When UNIQUE is specified, the page space is allocated a VSAM data space of its own. VSAM generates a name for the data space and builds a format-1 DSCB in the volume's VTOC to describe it. If SUBALLOCATION is specified, the name of the data space, not of the page space, appears in the VTOC. A data space must have been defined on the volume on which the page space is defined.

Auxiliary Storage Management recommends that UNIQUE be used for page space data sets. This will cause VSAM to allocate a single nonshared data space for the page data set. Space from another data space will not be used, eliminating one of the possibilities for multiple extent data sets which will cause Auxiliary Storage Management performance degradation.

Abbreviations: UNQ and SUBAL

UPDATEPW(password)

specifies an update-level password for the page space. Because the page space is a system data set, it cannot be opened or used by a user's program.

If a read password is the only password specified for the page space (that is, it is the highest-level password), it propagates upward and becomes the password for all higher levels. If you specify a higher-level password and do not specify an update password, the update password is null.

The passwords are cataloged in both the page space entry and its data component's entry. You should specify control-, update-, and read-level passwords to prevent the page space from being accessed if its volume is moved to a VS1 system. A VS2 system automatically prevents users from accessing page spaces and their data components.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, UPDATEPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, UPDATEPW('*CA' 'ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphanumeric or special character. For example, UPDATEPW(X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

Catalog (Page Space)

CATALOG(catname[/password])

specifies the name and password of the catalog in which the page space is to be defined.

When the CATALOG parameter identifies a user catalog, you must also supply a STEPCAT or JOBCAT DD statement to describe and allocate the user catalog. See "Order of Catalog Use: DEFINE" for information about the order in which catalogs are used.

catname

specifies the name of the catalog.

password

specifies a password. If the catalog is password protected, you must supply the catalog's update- or higher-level password. If no password is specified, VSAM asks the operator or TSO terminal user for the correct password.

Abbreviation: CAT

DEFINE PAGESPACE EXAMPLES

Define a Conventional (NOSWAP) Page Space: Example 1

In this example, a page space is defined.

```
//DEFPGSP1 JOB      ...
//STEP1   EXEC     PGM=IDCAMS
//VOLUME  DD       VOL=SER=VSER05,UNIT=2305,DISP=OLD
//SYSPRINT DD      SYSOUT=A
//SYSIN   DD       *
          DEFINE PAGESPACE -
              (NAME(SYS1.PAGE2) -
              CYLINDERS(10) -
              VOLUMES(VSER05) -
              CONTROLPW(PASSWD1) -
              UPDATEPW(PASSWD2) -
              READPW(PASSWD3))
/*
```

The job control language statements are:

- VOLUME DD, which describes the volume on which the data space is to be defined.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

Because no catalog is explicitly specified with a CATALOG parameter or with JOBCAT or STEPCAT DD statements, VSAM assumes that the data space and page space are to be defined in the master catalog.

The DEFINE PAGESPACE command defines a page space. This page space will have the UNIQUE attribute and hence a new data space will be created for its exclusive use. Its parameters are:

- NAME, which specifies the name of the page space, SYS1.PAGE2.
- CYLINDERS, which specifies that the page space is to occupy ten cylinders. The page spaces are never extended.
- VOLUMES, which specifies that the page space is to reside on volume VSER05.
- CONTROLPW, which specifies the control password for the page space is PASSWD1.
- UPDATEPW, which specifies the update password for the page space is PASSWD2.
- READPW, which specifies the read password for the page space is PASSWD3.

Because neither UNIQUE nor SUBALLOCATION has been specified, the page space defaults to UNIQUE. Because neither SWAP nor NOSWAP has been specified, the page space defaults to NOSWAP.

Define a SWAP Page Space: Example 2

In this example, a unique page space is defined. A unique page space differs from a suballocated page space in that a unique page space occupies an entire VSAM data space. When a unique page space is defined, a VSAM data space is automatically built for the page space. When space is allocated to a unique page space, the volume's VTOC

identifies the space with the name of the page space.

```
//DEFPGSP2 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DEFINE PAGESPACE -
  (NAME(SYS1.PAGE1) -
  CYLINDERS(10) -
  VOLUMES(VSER05) -
  SWAP -
  CONTROLPW(PASSWD1) -
  UPDATEPW(PASSWD2) -
  READPW(PASSWD3))
/*
```

The DEFINE PAGESPACE command defines a page space. Its parameters are:

- NAME, which specifies the name for the page space: SYS1.PAGE1. Because the page space is unique, the page space's name is put into the DSCB (in the volume's VTOC) that describes the space allocated to the page space.
- CYLINDERS, which specifies that the page space occupies 10 cylinders and cannot be extended.
- VOLUMES, which identifies the volume on which the page space is to reside. Because no DD statement describes the volume, the volume is dynamically allocated. Volume VSER05 must be mounted as permanently resident or reserved.
- SWAP, which specifies the page space will be used to store Local System Queue Area (LSQA) pages.
- CONTROLPW, which specifies the control password for the page space is PASSWD1.
- UPDATEPW, which specifies the update password for the page space is PASSWD2.
- READPW, which specifies the read password for the page space is PASSWD3.

Because neither UNIQUE nor SUBALLOCATION has been specified, the page space defaults to UNIQUE.

Because neither a CATALOG parameter nor a JOBCAT or STEPCAT DD statement was specified, VSAM assumes that the page space is to be defined in the master catalog.

DEFINE PATH

The format of the DEFINE command when it is used to define a path is:

DEFINE	PATH (NAME(entryname) PATHENTRY(entryname[/password]) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLPW(password)] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [TO(date) FOR(days)] [UPDATE NOUPDATE] [UPDATEPW(password)] [CATALOG(catname[/password])]
---------------	--

DEFINE PATH Parameters: Summary

The parameters of the DEFINE PATH command can be grouped as required and optional parameters.

The required parameters are:

- Entry type, which specifies that a path catalog entry is to be created (PATH).
- Name, which names the path and the entry that it points to (NAME and PATHENTRY).

The optional parameters are:

- Allocation, which identifies the catalog recovery volume that contains the alternate index or cluster that the path points to (FILE). This parameter is required when the path is defined in a recoverable catalog.
- Protection and integrity, which
 - identify the path to a console operator without revealing the path's entry name (CODE and ATTEMPTS).
 - identify a user-written-security verification routine for additional authorization verification (AUTHORIZATION).
 - specify a retention period for the path (FOR or TO).
 - associate passwords with the path (MASTERPW, CONTROLPW, UPDATEPW, and READPW).
 - identify the owner of the path (OWNER).
- Model, which identifies an existing path entry that is to be used as a model for the path being defined (MODEL).
- Update, which specifies whether the cluster's upgrade set is to be opened when the path (and its cluster) is opened (UPDATE or NOUPDATE).

- Catalog, which identifies the catalog that contains the alternate index or cluster entry that the path entry is to point to (CATALOG).

DEFINE may be abbreviated: DEF.

DEFINE PATH PARAMETERS

Required Parameters

PATH

specifies that a path is to be defined.

The PATH keyword is followed by parameters which name the path and its related object, and which specify the path's attributes.

NAME(entryname)

specifies the path's name.

entryname

can contain 1 to 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and special characters (the hyphen (-), and the 12-0 overpunch (X'C0')).

When the name contains more than eight characters, you divide the name into segments of 1 to 8 characters each and separate the segments with (.) periods. The name's first character (and each segment's first character) is either an alphabetic character or a national character. If you specify the path's name in the format entryname(modifier), only that portion of the path's name preceding the left parenthesis will be used. The modifier enclosed in parentheses will be ignored.

PATHENTRY(entryname[/password])

when the path consists of an alternate index and its base clusters, entryname identifies the alternate-index entry. When the path is opened to process data records, both the alternate index and the base cluster are opened.

When the path consists of a cluster without an alternate index, entryname identifies the cluster. You can define the path as though it were an alias for the cluster. This allows you to specify no-update access to the cluster, so that the upgrade set will not be required or updated when the cluster is opened. You can also establish protection attributes for the alternate name, separate from the protection attributes of the cluster.

If the cluster or alternate-index entry is password protected, you must supply the entry's master password. When you identify the catalog with the CATALOG parameter, you can supply the catalog's master password instead of the entry's password.

Abbreviation: PENT

Optional Parameters

ATTEMPTS(number|2)

specifies the maximum number of times the operator can try to enter a correct password in response to a prompting message.

This parameter can be coded, but only has effect when the path's master password is not null. A prompting message is issued only when the user has not already supplied the appropriate password.

number
is an integer from 0 to 7 and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Note to TSO users: At a TSO terminal, the logon password is checked first before the user is prompted to supply a password for the path. Checking the logon password counts as one attempt to obtain a password. If ATTEMPTS is not specified, the user has one attempt to supply the path's password because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])
specifies that a user-security-verification routine (USVR) is available for additional security verification. When a protected path is accessed and the user supplies a correct password other than the cluster's master password, the USVR receives control. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for details on the USVR.

If a USVR is loaded from an unauthorized library during Access Method Services processing, an abnormal termination will occur. See the section "Authorized Program Facility" in the chapter "Introduction."

This parameter can always be coded, but only has effect when the path's master password is not null.

entrypoint
specifies the name of the USVR. entrypoint can contain 1 to 8 alphanumeric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and special characters (the hyphen(-) and the 12-0 overpunch (X'C0')). The name's first character is either an alphabetic character or a national character.

string
specifies information to be passed to the USVR when it receives control to verify authorization. string can contain 1 to 255 EBCDIC characters.

When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single-quotation marks (for example, AUTHORIZATION(entrypoint, 'XDORIS*')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, 'XCA'ROLX')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint, SUE). The string can contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

CODE(code)

specifies a code name for the path. If an attempt is made to access a password-protected path without first supplying an appropriate password, a prompting message is issued to the operator's console. The prompting message includes the code name, which identifies the path without revealing its entry name.

This parameter can be coded, but only has effect when the path's master password is not null. A prompting message is issued only when the user hasn't already supplied the appropriate password. When CODE is not specified, the prompting message identifies the path with its entry name.

code

can contain 1 to 8 alphameric or special characters.

When the code contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the code in single-quotation marks (for example, CODE('*DORIS*')).

If the code contains a lowercase alphabetic character it is folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE('*CA''ROL*')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CODE(X'E2E4C5') is the same as CODE(SUE).

CONTROLPW(password)

specifies a control password for the path. Control interval processing is not permitted through a path; the control password will permit read and write operations against the base cluster.

If a read or update password is the only password specified for the object, it (the highest-level password) propagates upward and becomes the password for all higher unspecified levels.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

FILE(dname)
 specifies the name of a DD statement that identifies the recovery volume of the alternate index or cluster named with the PATHENTRY parameter. FILE is only required when the path is defined in a recoverable catalog.

The recovery volume is the first volume of the base cluster's index component when PATHENTRY names a key-sequenced cluster or an alternate index over a key-sequenced cluster.

The recovery volume is the first volume of the base cluster's data component when PATHENTRY names an entry-sequenced cluster or an alternate index over an entry-sequenced cluster.

When FILE is not specified, the cluster's recovery volume is dynamically allocated. The volume must be mounted as permanently resident or reserved.

MASTERPW(password)
 specifies a master password for the path. The master password allows all operations against the path.

If the master password is not specified, the path's highest-level password propagates upward and becomes the password for all higher levels, including the master password.

If all passwords are null, ATTEMPTS, AUTHORIZATION, and CODE can be coded but have no effect until the master password is specified.

password
 is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation: MRPW

**MODEL(entryname[/password]
 [b catname[/password]])**
 identifies an existing path entry that is to be used as a model for the path being defined. For details about how a model is used, see "How to Use One Object as a Model for Another Object."

You can use some attributes of the model and override others by explicitly specifying them in the definition of the path. When you do not want to add or change any attributes, you specify only the entry-type (PATH), the path's name, its alternate-index's (or cluster's) name, and

the model-entry's name.

entryname

names the entry to be used as a model. entryname must name a path entry.

password

specifies a password. If the model entry is password protected and it is cataloged in a password-protected catalog, you must supply the read (or higher level) password of either the model entry or its catalog. If you specify both passwords, the catalog's password is used.

If you are not specifying new protection attributes for the path (that is, the model's passwords and protection attributes are being copied), you must supply the master password of either the model entry or its catalog.

catname

names the model-entry's catalog. You must identify the catalog that contains the model entry for either of these cases:

- You specify the catalog's password instead of the model-entry's password
- The model-entry's catalog is not identified with a JOBCAT or STEPCAT DD statement, and is not the master catalog.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved. See "Order of Catalog Use: DEFINE" for information about the order in which a catalog is selected when the catalog's name is not specified.

OWNER(ownerid)

specifies the identification of the path's owner.

ownerid

can contain 1 to 8 EBCDIC characters.

When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('DORIS')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('CA'ROL')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

Note to TSO users: If the owner is not identified with the OWNER parameter, the TSO user's userid becomes the ownerid.

READPW(password)

specifies a read password for the path. The read password permits read operations against the base cluster's data records.

password

is a 1 to 8 alphanumeric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('*CA''ROLX*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

TO(date)|FOR(days)

specifies the retention period for the path. The path is not automatically deleted when the expiration date is reached. When a retention period is not specified, the path can be deleted at any time.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the Julian date (001, for January 1, through 365, for December 31), through which the path is to be kept before it is allowed to be deleted.

FOR(days)

specifies the number of days for which the entry is to be kept before it is allowed to be deleted. days can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. The maximum number that can be specified for days is 9999. If the number specified is 0 through 1830, the retention period is the number of days specified. If the number specified is between 1831 and 9999, the retention period is through the year 1999.

UPDATE|NOUPDATE

specifies whether the base-cluster's upgrade set is to be allocated when the path is opened for processing.

The upgrade set is a group of alternate indexes associated with the base cluster. The alternate indexes are opened whenever the base cluster is opened. When records in the base cluster are modified or deleted, and when records are added to the base cluster, each alternate index in the base-cluster's upgrade set is modified to reflect the change in the cluster's data, just as a key-sequenced cluster's index is modified each time the cluster's data changes.

When a path points to a base cluster that has a large upgrade set (that is, many alternate indexes are associated with the base cluster), and the path is defined with the NOUPDATE attribute, you can open the path, and consequently the base cluster, and none of the alternate indexes will be opened.

NOUPDATE specifies that, when opening the path, the path's base cluster is to be allocated and the base cluster's upgrade set is not to be allocated. You can specify the NOUPDATE attribute for the path even though the UPGRADE attribute is set for one of the base-cluster's alternate indexes.

Abbreviations: UPD and NUPD

UPDATEPW(password)

specifies the update password for the path. The update password permits read and write operations against the base cluster's data records.

If a read password is the only password specified for the object (that is, it is the highest-level password), it propagates upward and becomes the password for all higher levels. If you specify a higher-level password and don't specify an update password, the update password is null.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, UPDATEPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, UPDATEPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, UPDATEPW(X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

Catalog (Path)

The catalog parameter allows you to supply the name and password of the catalog in which the path is to be defined.

CATALOG(catname[/password])

identifies the catalog that contains the entry of the cluster or alternate index named in the PATHENTRY parameter.

See "Order of Catalog Use: DEFINE" for information about the order in which a catalog is selected if the catalog's name is not specified.

If the cluster's or alternate-index's entry is password protected and its catalog is also password protected, you must specify the master password for either the entry or the catalog.

catname

specifies the catalog's name.

password

specifies the catalog's master password.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved.

Abbreviation: CAT

DEFINE PATH EXAMPLE

Define a Path: Example 1

In this example, a path is defined. Previous examples illustrated the definition of the path's alternate index, EXAMPLE.AIX, and the alternate index's base cluster, EXAMPLE.KSDS2. The alternate index, path, and base cluster are defined in the same catalog, AMASTCAT.

```
//DEFPATH JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
        DEFINE PATH -
                (NAME(EXAMPLE.PATH) -
                 PATHENTRY(EXAMPLE.AIX/AIXMRPW) -
                 READPW(PATHRDPW) ) -
                CATALOG(AMASTCAT/MCATUPPW)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE PATH command builds a path entry to define the path EXAMPLE.PATH. A copy of the path entry is placed in the catalog recovery area, since the path entry is being defined into a recoverable catalog. The catalog recovery area on volume VSER04 is dynamically allocated. The command's parameters are:

- NAME, which specifies that the path's name is EXAMPLE.PATH.
- PATHENTRY, which identifies the alternate index, EXAMPLE.AIX, that the path provides access to.
- READPW, which specifies the path's read password, PATHRDPW.
- CATALOG, which in this example supplies the master catalog's update password, MCATUPPW.

DEFINE SPACE

The format of the DEFINE command when it is used to define a data space is:

DEFINE	SPACE ([CANDIDATE CYLINDERS(primary[b secondary]) RECORDS(primary[b secondary]) b RECORDSIZE(average b maximum) TRACKS(primary[b secondary]) VOLUMES(volser[b volser...]) [FILE(dname)])] [CATALOG(catname[/password])])
--------	--

The parameters of the DEFINE SPACE command can be divided into required and optional parameters.

The required parameters are:

- Allocation, which specifies the amount and type of space to be allocated to the data space, and the volume that is to contain the data space (VOLUMES, CYLINDERS, RECORDS, TRACKS, RECORDSIZE, and CANDIDATE).

The optional parameters are:

- Allocation, which:
 - Specifies the DD statement which identifies and allocates the volume where the space is to be allocated (FILE).
- Catalog, which identifies the catalog in which the data space is to be defined (CATALOG).

DEFINE may be abbreviated: DEF.

DEFINE SPACE PARAMETERS

Required Parameters

SPACE

specifies that a data space is to be defined. You can also use the DEFINE SPACE command to reserve a volume for VSAM's future use. The SPACE keyword is followed by the allocation parameters, enclosed in parentheses, that describe the space, and by the catalog parameter that identifies the catalog that owns the volume.

Abbreviation: SPC

CANDIDATE|

CYLINDERS(primary[b secondary])|
RECORDS(primary[b secondary])b RECORDSIZE(average b maximum)|
TRACKS(primary[b secondary])

specifies a volume(s) is to be reserved for VSAM's future use or specifies the amount of space to be allocated in terms of cylinders, number of records, or tracks.

CANDIDATE

specifies that the volumes listed in the VOLUMES parameter are reserved for future use by VSAM. An ownership indicator is set in a format-4 DSCB in the volume's table of contents (VTOC). The volumes are reserved for use by the catalog, but no space is allocated (other than the primary CRA allocation if the catalog is recoverable). Each volume so reserved is called a candidate volume. If CANDIDATE is not

specified, you must specify one of the following parameters: CYLINDERS, TRACKS, or RECORDS.

If a VSAM data space already exists on the volume, you cannot specify CANDIDATE. (You can, however, use the ALTER command to identify the volume as a candidate volume for a VSAM object other than the volume's catalog.)

Abbreviation: CAN

CYLINDERS(primary[b secondary])
RECORDS(primary[b secondary])
b RECORDSIZE(average b maximum)
TRACKS(primary[b secondary])

specifies the amount of space to be allocated in terms of tracks, cylinders, or number of records. If RECORDS is specified, the space required is calculated in terms of the number of records, but the space is allocated by tracks. However, if you specify TRACKS or RECORDS and if the minimum number of tracks should exceed a cylinder, space is allocated in terms of cylinders.

primary

specifies the initial amount of space (primary extent) to be allocated to the data space. primary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

secondary

specifies the amount of space that is to be allocated to the data space each time it is extended. Once the primary extent is filled, the data space can expand to include a maximum of 15 secondary extents if you have specified a secondary-extent amount. secondary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

The secondary space allocation is rounded upward to the nearest whole cylinder regardless of your specification in terms of TRACKS or RECORDS. Consequently, the total amount of space allocated to the data space might exceed the predicted limit.

When the data space is the first data space on a volume that belongs to a recoverable catalog, the primary amount is increased by one cylinder.

You must specify one, and only one, of the following parameters: CANDIDATE, CYLINDERS, TRACKS, or RECORDS. When you specify RECORDS, you must also specify RECORDSIZE. When you specify CYLINDERS or TRACKS, you cannot specify RECORDSIZE.

Abbreviations: CYL, REC, and TRK

RECORDSIZE(average b maximum)

specifies the average and maximum lengths, in bytes, of the records.

average and maximum can be any integer value between 1 and 32,761, expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. When you specify the RECORDS parameter, you must also specify RECORDSIZE. Otherwise, you cannot specify RECORDSIZE.

Abbreviation: RECSZ

VOLUMES(volser[b volser...])

specifies the volumes on which data spaces are to be defined. If two or more volumes are specified and an amount of space is specified (that is, TRACKS, CYLINDERS, or RECORDS and RECORDSIZE is specified) the amount specified for the primary allocation is allocated on each volume. The specified volumes must all be of the same device type. When

the volume already has a data space on it, subsequent data spaces can be defined and must be cataloged in the catalog that owns the volume. A data space can be defined on a mass storage volume.

A volume serial number, volser, may contain 1 to 6 alphanumeric, national (@, #, and \$), and, except under TSO, special characters. The special characters include commas, blanks, semicolons, parentheses, slashes, asterisks, periods, quotation marks, ampersands, plus signs, hyphens, and equal signs. A volume serial number must be enclosed in single quotation marks if it contains a special character. Single quotation marks within a volume serial number must be coded as two single quotation marks.

Abbreviation: VOL

Optional Parameters

FILE(dname)

specifies the name of the DD statement that identifies the device type and volumes to be used for space allocation. You cannot use concatenated DD statements to describe more than one volume.

If FILE is not specified and the volume is physically mounted, the volume(s) identified with the VOLUMES parameter is dynamically allocated. Each dynamically allocated volume must be mounted as permanently resident or reserved.

Catalog (Space)

CATALOG(catname[/password])

identifies the catalog in which the data space is to be defined. The catalog parameter is used to supply the name and password, when required, of the catalog that contains the volume entry that describes the volume on which the data space is to be allocated. If the DEFINE SPACE command is used to allocate the volume's first VSAM data space or to identify the volume as a candidate volume, the catalog parameter identifies the catalog that is to contain the volume's catalog entry and is to own the volume. (For more information about VSAM volume ownership, see the section "VSAM Volume Ownership" in the "Introduction.")

If the catalog's volume is physically mounted, it is dynamically allocated. See "Order of Catalog Use: DEFINE" for information about the order in which catalogs are used.

catname

specifies the name of the catalog.

password

specifies a password. If the catalog is password protected, you must supply the update- or higher-level password.

Abbreviation: CAT

DEFINE SPACE EXAMPLE

Define a Data Space: Example 1

In this example, a data space is defined. The data space is to be used to allocate space for VSAM clusters that are subsequently defined.

```
//DEFSPC1 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE SPACE -
          (CYLINDERS(100 10) -
          VOLUMES(VSER05))
/*
```

VSAM assigns a name to the data space and keeps track of it for future data set space allocation. Because no user catalog is specified in this example, the master catalog is assumed as the default. (It is the first catalog found because no JOBCAT or STEPCAT DD statements were specified. See the previous section "Order of Catalog Use: DEFINE" for more details.)

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE SPACE command defines a VSAM data space. Its parameters are:

- CYLINDERS, which specifies that 100 cylinders are to be allocated for the data space. When the data space is extended, it is to be extended in increments of ten cylinders.
- VOLUMES, which specifies the volume serial number of the volume on which the data space is to be defined: VSER05.

Access Method Services defines the data space and allocates its space on volume VSER05. The data space is cataloged in the master catalog, because the DEFINE command does not include a CATALOG parameter, and because no JOBCAT or STEPCAT DD statements are specified. All future data spaces and clusters on volume VSER05 must also be cataloged in the master catalog. NonVSAM data sets on the volume reside in areas that are not allocated to a VSAM data space.

To reserve a 3330 volume for VSAM's exclusive use, define the data space on it as 403 cylinders (this allows one cylinder to be used for the volume's table of contents—VTOC—and the volume's label) or define other data spaces on the volume with additional DEFINE SPACE commands (which can be given in the same job step) so that all the volume's space is part of a VSAM data space.

DEFINE USERCATALOG

When you define a catalog, you can specify attributes for the catalog as a whole and for the catalog's components.

The format of the DEFINE command when it is used to define a catalog is:

DEFINE	<pre>USERCATALOG MASTERCATALOG (NAME(entryname) [CYLINDERS(primary[b secondary]] RECORDS(primary[b secondary]] TRACKS(primary[b secondary])) VOLUME(volser) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size 3072)] [CODE(code)] [CONTROLPW(password)] [DESTAGWAIT NODESTAGWAIT] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [RECOVERABLE NOTRECOVERABLE] [TO(date) FOR(days)] [UPDATEPW(password)] [WRITECHECK NOWRITECHECK]) [DATA ([BUFFERSPACE(size)] [CYLINDERS(primary[b secondary]] RECORDS(primary[b secondary]] TRACKS(primary[b secondary]))] [DESTAGWAIT NODESTAGWAIT] [RECOVERABLE NOTRECOVERABLE] [WRITECHECK NOWRITECHECK]) [INDEX ([CYLINDERS(primary)] RECORDS(primary)] TRACKS(primary)] [DESTAGWAIT NODESTAGWAIT] [WRITECHECK NOWRITECHECK]) [CATALOG(mastercatname[/password])]</pre>
--------	---

DEFINE USERCATALOG Parameters: Summary

The parameters of the DEFINE USERCATALOG|MASTERCATALOG command can be divided into required and optional parameters.

The required parameters are:

- Entry type, which specifies that a user catalog is to be defined (USERCATALOG or MASTERCATALOG).
- Name, which names the catalog (NAME).
- Allocation, which
 - Specifies the space to be allocated for the catalog's data space (one of CYLINDERS, RECORDS, or TRACKS).
 - Specifies the volume on which the catalog is to reside (VOLUME).

The optional parameters are:

- Allocation, which describes:
 - The name of the DD statement which describes the catalog volume (FILE).
 - The space for the catalog itself (one of CYLINDERS, RECORDS, or TRACKS).
 - The space to be used for buffers when the catalog is accessed (BUFFERSPACE).
 - Whether a catalog recovery area is to be defined (RECOVERABLE or NOTRECOVERABLE).
- Protection and integrity, which
 - Associate passwords with the catalog. (MASTERPW, CONTROLPW, UPDATEPW, READPW).
 - Provide a mechanism by which the console operator can be prompted to supply a password without disclosing the name of the catalog (CODE).
 - Identify a user-written routine for additional authorization verification (AUTHORIZATION).
 - Identify the owner of the catalog (OWNER).
 - Specify the time for which the catalog is to be kept (TO or FOR).
 - Specify whether the write-check operation is to be performed (WRITECHECK or NOWRITECHECK).
 - Specify whether a user catalog that is stored on a mass storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it (DESTAGEWAIT or NODESTAGEWAIT).
- Model, which identifies an existing catalog from which attributes are to be copied for the definition of a new user catalog (MODEL).
- Components, which allows you to specify attributes separately for the catalog's data and index components (DATA and INDEX).
- Catalog, which supplies the name and password, when required, of the master catalog (CATALOG).

DEFINE can be abbreviated: DEF.

DEFINE USERCATALOG PARAMETERS

Required Parameters

USERCATALOG(options)

specifies that a user catalog is to be defined. USERCATALOG is followed by the parameters specified for the catalog as a whole.

(In OS/VS2 MVS, when you specify MASTERCATALOG the result is the same; a user catalog is created. You can, however, establish a user catalog as a master catalog at IPL time. See "Appendix H: Changing the Attributes of the VSAM Master Catalog" for a description of this procedure.)

Abbreviation: UCAT

NAME(entryname)

specifies the name of the catalog being defined.

The name may contain from 1 through 44 alphameric characters (A through Z, and 0 through 9), national characters (@, #, and \$), and two special characters (the hyphen and the 12-0 overpunch).

Names containing more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of any name or name segment must be either an alphabetic or a national character. If you specify the name in the format entryname(modifier), only that portion of the name preceding the left parenthesis will be used. The modifier enclosed in parentheses will be ignored.

CYLINDERS(primary[b secondary])**RECORDS(primary[b secondary])****TRACKS(primary[b secondary])**

specifies the amount of space to be allocated in terms of tracks, cylinders, or number of records. You can specify the amount of space as a parameter of USERCATALOG, as a parameter of USERCATALOG and DATA, or as a parameter of USERCATALOG, DATA, and INDEX. See the section "How Space is Assigned to a Catalog" for a description of the difference in space allocation depending on which combination you use. See the section "Estimating the Catalog's Space Requirements" for information about estimating the amount of space to be specified for a catalog.

When you specify less than one cylinder of space, and the catalog is defined as recoverable (see the RECOVERABLE parameter), your DEFINE command will fail.

primary[b secondary]

specify the size of the primary and secondary extents to be allocated. Once the primary extent is filled, the space can expand to include a maximum of 13 additional secondary extents if you have specified a secondary allocation amount.

primary and secondary can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. Secondary allocation should be specified in case the catalog has to be extended.

Abbreviations: CYL, REC, and TRK

VOLUME(volser)

specifies the volume that is to contain the catalog. The volume cannot be currently owned by any other VSAM catalog. A user catalog can be defined on a mass storage volume.

A volume serial number, volser, may contain 1 to 6 alphameric, national (@, #, and \$), and special characters; the special characters include commas, blanks, semicolons, parentheses, slashes, asterisks, periods, quotation marks, ampersands, plus signs, hyphens, and equal signs. A volume serial number must be enclosed in single quotation marks if it contains a special character. Single quotation marks within a volume serial number must be coded as two single quotation marks.

The VOLUME parameter interacts with other DEFINE CATALOG parameters. You should take care to ensure that the volume you specify for the catalog can satisfy the catalog's other attributes:

- **CYLINDERS, RECORDS, TRACKS:** The volume contains enough unallocated space to satisfy the catalog's primary space requirement. (Space on the volume might already be allocated to nonVSAM data sets and system data sets. However, the volume cannot be owned by another VSAM

catalog—that is, no space on the volume can already be allocated to VSAM objects.)

- FILE: The volume information supplied with the DD statement is consistent with the information specified for the catalog and its components.

Abbreviation: VOL

Optional Parameters

ATTEMPTS(number|2)

specifies the maximum number of times the operator can try to enter a correct password in response to a prompting message.

number

is an integer from 0 to 7 and can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form. If 0 is specified, the operator is not prompted and is not allowed to enter a password from the console.

Note to TSO Users: At a TSO Terminal the logon password is checked first before the user is prompted to supply a password for the catalog. Checking the logon password counts as one attempt to obtain a password. If ATTEMPTS is not specified, the user has one attempt to supply the catalog's password, because the default is 2.

Abbreviation: ATT

AUTHORIZATION(entrypoint[b string])

specifies that a user-security-verification routine (USVR) is available for additional security verification. When a protected catalog is accessed and the user supplies a correct password other than the catalog's master password, the USVR receives control. See "User-Written Exit Routines" in OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for information on the user-security-verification routine.

entrypoint

specifies the name of the USVR. entrypoint can contain 1 to 8 alphanumeric characters (A through Z, and 0 through 9), or national characters (@, #, \$), and special characters (the hyphen (-) and the 12-0 overpunch (X'CO')). The name's first character is either an alphabetic character or a national character.

string

specifies information to be passed on to the USVR when it receives control to verify authorization. string can contain 1 to 255 EBCDIC characters.

When the string contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the string in single-quotation marks (for example, AUTHORIZATION(entrypoint, '*DORIS*')).

When the string contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, AUTHORIZATION(entrypoint, '*CA''ROL*')).

You can code string in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, AUTHORIZATION(entrypoint, X'E2E4C5') is the same as AUTHORIZATION(entrypoint, SUE). The string

can contain up to 255 hexadecimal characters when expressed in this form, resulting in up to 128 bytes of information.

Abbreviation: AUTH

BUFFERSPACE(size|3072)

specifies the minimum space, in bytes, to be provided for buffers when using the catalog. Decimal values you can specify are 3072, 4096, 5120, 6144, 7168, and 8192.

size

is the amount of space, in bytes, to be provided for buffers. size can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

Abbreviation: BUFSP

CODE(code)

specifies a code name for the catalog being defined. If an attempt is made to access a password-protected catalog without a password, the code name is used in a prompting message; the code enables the operator to be prompted for the password without disclosing the name of the catalog. If CODE is not specified, the operator is prompted with the name of the catalog.

code

can contain 1 to 8 alphameric or special characters. When the code contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the code in single-quotation marks (for example, CODE('*DORIS*')).

If the code contains a lowercase alphabetic character it will be folded to an uppercase alphabetic character.

When the code contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CODE('*CA''ROLX*')).

You can code code in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CODE (X'E2E4C5') is the same as CODE(SUE).

CONTROLPW(password)

specifies a control password for the catalog being defined. The control password permits the same operations as the update password.

If a read or update password is the only password specified for the catalog, it (the highest-level password) propagates upward and becomes the password for all higher levels. When the master password is specified but the control password is not specified, the control password is null.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example,

CONTROLPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, CONTROLPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, CONTROLPW(X'E2E4C5') is the same as CONTROLPW(SUE).

Abbreviation: CTLPW

DESTAGEWAIT|NODESTAGEWAIT

specifies whether a user catalog that is stored on a mass storage volume is to be destaged synchronously or asynchronously with respect to the program that closes it.

DESTAGEWAIT

indicates that destaging is to be completed before control is returned from VSAM to the program that issues the CLOSE macro. VSAM can thus notify the program whether destaging was successful.

Abbreviation: DSTGW

NODESTAGEWAIT

indicates that notification of unsuccessful destaging is to be made only by a message to the operator.

Abbreviation: NDSTGW

If the user catalog is not stored on a mass storage volume, the attribute is ineffective until the direct access storage volume the catalog is stored on is converted to a mass storage volume (by way of the CONVERTV command, which is described in OS/VS Mass Storage System (MSS) Services: Reference Information).

If one of these parameters is specified for the data component of a catalog and the other parameter is specified for the index component, the parameter specified for the index component applies to both components, because the sequence set of the index of a catalog is imbedded in the data.

FILE(dname)

specifies the name of the DD statement that identifies the device and volume to be used for the catalog. The DD statement should specify DISP=OLD to prevent premature space allocation on the volume. If FILE is not specified and the catalog's volume is physically mounted, the volume identified with the VOLUME parameter is dynamically allocated. The volume must be mounted as permanently resident or reserved.

MASTERPW(password)

specifies a master password for the catalog being defined. A catalog must be master password protected in order to password protect any VSAM data sets cataloged in it. The AUTHORIZATION, CODE, and ATTEMPTS parameters have no effect unless the catalog has a master password associated with it. The master password allows all operations; it is required to open the catalog as a data set.

If a master password is not specified for the catalog, the catalog's highest-level password propagates upward and becomes the password for all higher levels, including the master password. The master password is null when no other passwords are specified.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, MASTERPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, MASTERPW('*CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, MASTERPW(X'E2E4C5') is the same as MASTERPW(SUE).

Abbreviation; MRPW

MODEL(entryname[/password][b catname[/password]]) specifies that an existing master or user catalog is to be used as a model for the user catalog being defined. For details about how a model is used, see "How to Use One Object as a Model for Another Object."

It is possible to use an already defined master or user catalog as a model for another user catalog. When one entry is used as a model for another, its attributes are copied as the new entry is defined. You may use some attributes of the model and override others by explicitly specifying them in the definition of the user catalog.

If a model is used, you must specify certain parameters even though no attributes are to be changed or added. The name of the user catalog to be defined and volume and space information always must be specified as parameters of USERCATALOG.

entryname

specifies the name of the master or user catalog to be used as a model.

password

specifies a password. If the catalog to be used as a model is password protected, a password is required. If you specify both passwords (that is, the password following entryname and the password following catname), the password following catname is used for authorization. If the protection attributes are to be copied, enter the master password of the catalog being used as a model. If passwords are not to be copied, any password except the master can be used.

catname

specifies the name of the catalog to be used as a model. This parameter is required if the model catalog is neither the master catalog nor a catalog identified by a JOBCAT or STEPCAT DD statement.

OWNER(ownerid)

specifies the identification of the owner of the catalog being defined.

ownerid

can contain 1 to 8 EBCDIC characters.

When the ownerid contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the ownerid in single-quotation marks (for example, OWNER('*DORIS*')).

When the ownerid contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, OWNER('*CA''ROL*')).

You can code ownerid in hexadecimal form, where two hexadecimal characters represent one EBCDIC character. For example, OWNER(X'E2E4C5') is the same as OWNER(SUE).

READPW(password)

specifies a read password for the catalog being defined. The read password permits the user to list the catalog's entries (passwords and protection attributes are listed only when the master-level password is supplied).

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, READPW('*DORIS*')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, READPW('CA''ROL*')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, READPW(X'E2E4C5') is the same as READPW(SUE).

Abbreviation: RDPW

RECOVERABLE|NOTRECOVERABLE

specifies that a catalog recovery space is to be created on each volume owned by the catalog.

On the catalog's volume, the catalog recovery space is allocated from the catalog's data space. On subsequent volumes owned by the catalog, the catalog recovery space is allocated from the first data space defined on the volume. Primary allocation for a CRA is one cylinder. Once a primary extent is filled, the CRA can expand to include a maximum of 15 additional extents. Allocation for each secondary extent is one cylinder. The number of records a CRA can contain varies with the device type (for example, for a 3330 with 20 control intervals per track, each extent will accommodate approximately 370 records).

Abbreviations: RVBL and NRVBL

TO(date)|FOR(days)

specifies the retention period for the catalog being defined. If no value is coded, the catalog can be deleted whenever it is empty.

TO(date)

specifies the date, in the form yyddd, where yy is the year and ddd is the number (001 through 365) of the day, through which the catalog being defined is to be kept.

FOR(days)

specifies the number of days for which the catalog being defined is to be kept. The maximum number that can be specified is 9999. If the number specified is 0 through 1830, the catalog is retained for the number of days specified; if the number is between 1831 and 9999, the catalog is retained through the year 1999.

days

can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form.

UPDATEPW(password)

specifies an update password for the catalog being defined. The update password permits entries to be added to the catalog being defined. The catalog must be update password protected in order to prevent unauthorized users from cataloging data sets.

When you anticipate using the catalog to contain many nonVSAM entries, you should consider allowing the catalog's update password to be null (that is, to specify a master password but not an update password). A null update password allows you to alter and delete nonVSAM entries without having to supply the catalog's password.

The phrase "an update- or higher-level password" means:

- If the update password exists, you are required to correctly supply either the update, control, or master password before Access Method Services processing continues.
- If the update password is null (that is, it does not exist), no prompting message is issued and no password is required.

You can protect the catalog and its VSAM entries by specifying a master password. You can allow the unrestricted alteration and deletion of nonVSAM entries and the addition of all types of entries by not specifying an update password.

If a read password is the only password specified for the catalog (that is, it is the highest-level password), it propagates upward and becomes the password for all higher levels. If you specify a higher-level password and do not specify an update or read password, the update password is null.

password

is a 1 to 8 alphameric or special character password.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the password contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the password in single-quotation marks (for example, UPDATEPW('XDORIS')).

When the password contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks

(for example, UPDATEPW('XCA'ROLX')).

You can code password in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, UPDATEPW(X'E2E4C5') is the same as UPDATEPW(SUE).

Abbreviation: UPDPW

WRITECHECK|NOWRITECHECK

specifies whether the catalog is to be checked by a direct-access device operation called write-check when a record is written to the device. If WRITECHECK is specified, a record is written and then read, without data transfer, to test for the data check condition.

Abbreviations: WCK and NWCK

Components (Catalog)

Attributes can be specified separately for the catalog's data and index components.

DATA(optional parameters)

specifies the data component's attributes. If an optional parameter is not specified as a parameter of DATA, the data component shares the value or attribute specified for the catalog as a whole.

Except for space allocation, attributes specified as parameters of DATA override the same attributes specified as parameters of USERCATALOG. Attributes specified as parameters of DATA also override the same attributes of a model catalog's data component.

Parameters that can be specified for the catalog's data component are:

- **BUFFERSPACE:** Specifies the amount of buffer space that the user's program is to provide from its address space when the program opens the data component as a separate VSAM object.
- **CYLINDERS, RECORDS, or TRACKS:** You can specify an amount of space to be allocated to the data component. You must also specify an amount of space for the catalog's data space. See the section "How Space is Assigned to a Catalog" for a description of the significance of specifying space allocation as a parameter of DATA.
- **RECOVERABLE or NOTRECOVERABLE:** Specifies whether a catalog recovery area is to be created on each volume owned by the catalog.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the catalog's records are to be checked by the direct-access device when they are written into the data component.
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the catalog's data component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the program that closes it.

INDEX(optional parameters)

specifies the index component's attributes. If an optional parameter is not specified as a parameter of INDEX, the index component shares the value or attribute specified for the catalog as a whole.

Except for space allocation, attributes specified as parameters of INDEX override the same attributes specified as parameters of USERCATALOG. Attributes specified as parameters of INDEX also override the same attributes of a model catalog's index component.

Parameters that can be specified for the catalog's index component are:

- **CYLINDERS, RECORDS, or TRACKS:** You can specify an amount of space to be allocated to the index component. When you specify one of the space-quantity parameters as a parameter of INDEX, you must also specify a space-quantity parameter as a parameter of DATA. TRACKS, CYLINDERS, or RECORDS must have been specified as a parameter of DATA to be specified as a parameter of INDEX. See the section "How Space is Assigned to a Catalog" for a description of the significance of specifying space allocation as a parameter of DATA and INDEX.
- **DESTAGEWAIT or NODESTAGEWAIT:** Specifies whether the catalog's index component, stored on a mass storage volume, is to be destaged synchronously or asynchronously with respect to the program that closes it.
- **WRITECHECK or NOWRITECHECK:** Specifies whether the index records are to be checked by the direct-access device when they are written into the index.

Catalog (Catalog)

The catalog parameter is used to supply the name and password, when required, of the master catalog when a user catalog is to be defined.

CATALOG(mastercatname[/password])

specifies the name and password of the master catalog. If the master catalog is password protected, its update- or higher-level password must be provided in this parameter or in response to prompting.

mastercatname

is the name of the master catalog that is required when a user catalog is being defined.

password

specifies the update- (or higher-level) password of the master catalog that is required when a user catalog is being defined, if the master catalog is password protected.

Abbreviation: CAT

DEFINE USERCATALOG EXAMPLES

Defining a User Catalog and Determining the Catalog's Space Requirements:
Example 1

In this example, a user catalog is defined. The user catalog's data space is not available to contain other VSAM objects. The catalog is to be large enough to contain information about:

- 100 key-sequenced clusters (that is, indexed VSAM data sets)
- 10 entry-sequenced clusters (that is, sequential nonindexed VSAM data sets)
- 5 key-sequenced clusters with alternate indexes to be upgraded
- 5 entry-sequenced clusters with alternate indexes to be upgraded
- 10 alternate indexes
- 10 paths
- 200 aliases
- 10 generation data sets
- 1000 nonVSAM data sets.
- 5 volumes (the number to be controlled by the catalog). Each volume is a 3330 volume and contains 20 data spaces.

To determine how much space, in records and tracks, is required for the primary allocation of a to-be-defined catalog in your system, see the section, "Estimating the Catalog's Space Requirements." The number of records for the primary allocation of the catalog in this example is shown in the completed worksheet in Figure 9.

When you define the catalog's space parameters, use the following format:

```
DEFINE USERCATALOG -
      (TRACKS(prim secn) ... )
```

where:

- prim = the minimum amount of space for the catalog data-space's primary allocation. In this example, all of the data space is assigned to the catalog itself. The calculated value of prim should be rounded upward to the next whole number of tracks which is an even multiple of the catalog allocation unit for the device type being used (see the section "How Space is Assigned to a Catalog" for the size of an allocation unit).

Because the catalog resides on a 3330 volume, the value of prim is:

$$\text{prim} = .09N + 3 = 152.4 = 153 \text{ tracks}$$

Variable Quantities	Formulas	Estimates
Basic requirement = 10 records		10
A = number of key-sequenced clusters	Ax3	300
A ¹ = number of key-sequenced clusters with alternate indexes to be upgraded	A ¹ x4	20
B = number of entry-sequenced clusters	Bx2	20
B ¹ = number of entry-sequenced clusters with alternate indexes to be upgraded	B ¹ x3	15
C = number of relative-record clusters	Cx2	0
D = number of alternate indexes	Dx3	30
E = number of path entries	E	10
F = number of nonVSAM data set entries	F	1000
G = number of generation data group entries	G	10
H = number of alias entries	H	200
I = number of page spaces	I x2	0
J = number of volumes, depending on device type, owned by the catalog:		
J ¹ = number of 2305 volumes	J ¹ x2	0
J ² = number of 2314/2319 volumes	J ² x3	20
J ³ = number of 3330 and 3340/3344 volumes	J ³ x4	0
J ⁴ = number of 3330 Model 11 and 3350 volumes	J ⁴ x6	
K = for each key-sequenced cluster and alternate index (KSDS) with space on more than two volumes, add "1" for each additional group of one to five volumes:		
K ¹ = number of KSDSs with 3 to 7 volumes	K ¹	0
K ² = number of KSDSs with 8 to 12 volumes	K ² x2	0
K ³ = number of KSDSs with 13 to 17 volumes	K ³ x3	0
L = for each entry-sequenced cluster and relative-record cluster (ESDS) with space on more than five volumes, add "1" for each additional group of one to eight volumes:		
L ¹ = number of ESDSs with 6 to 13 volumes	L ¹	0
L ² = number of ESDSs with 14 to 21 volumes	L ² x2	0
M = for each group of four data spaces on a volume, add "1"	M	25
N = number of entry records required for the catalog's data component (total of above)	N	1660

Figure 9. Completed Worksheet for Determining a Catalog's Space Requirements

- `secn` = the minimum amount of space for each secondary extent of the catalog's data space.

The number of tracks required for the catalog's primary extent is at least 153. This figure is an even multiple of an allocation unit for a 3330 device (3); therefore, no further rounding is necessary. If VSAM extends the catalog, it should extend it in increments of 18 tracks.

```
//DEFCAT1 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE USERCATALOG( -
            NAME(MYCAT) -
            VOLUME(VSER04) -
            NOTRECOVERABLE -
            TRACKS(153 18) ) -
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE USERCATALOG command defines a user catalog. Its parameters are:

- NAME, which specifies the name of the catalog, MYCAT.
- VOLUME, which specifies that the catalog is to reside on the volume whose serial number is VSER04.
- NOTRECOVERABLE, which specifies that no recovery space is to be allocated for this catalog.
- TRACKS, which specifies that the catalog's data space is to be 153 tracks. When the user catalog is extended, VSAM extends it in increments of 18 tracks.
- Values and attributes that apply by default to the catalog are:
 - BUFFERSPACE = 3072 bytes
 - ATTEMPTS = 2
 - NOWRITECHECK
 - NODESTAGEWAIT
- Null attributes and values of the catalog are:
 - All passwords
 - CODE
 - AUTHORIZATION
 - TO and FOR
 - OWNER

Define a User Catalog Leaving Available Suballocatable Space: Example 2

In this example, a user catalog is defined. The catalog will actually occupy only a portion of the space defined, leaving the remainder of the space available for the suballocation of other VSAM data sets.

```

//DEFCAT2 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE USERCATALOG( -
              NAME(D27UCAT1) -
              MASTERPW(MRPWD27) -
              UPDATEPW(UPPWD27) -
              FOR(365) -
              CYLINDERS(150 5)) -
              VOLUME(VSER02) ) -
          DATA( -
              CYLINDERS(8 5) -
          INDEX( -
              CYLINDERS(5) ) -
          CATALOG(AMASTCAT/MRCATPW2)
/*
    
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE USERCATALOG command defines a user catalog, D27UCAT1. Its parameters are:

- NAME, which names the user catalog, D27UCAT1.
- MASTERPW and UPDATEPW, which specify master and update passwords for the catalog: MRPWD27 and UPPWD27.
- FOR, which specifies that the user catalog is to be retained for 365 days.
- CYLINDERS, which specifies that 150 cylinders are to be allocated for the user catalog's data space. When the user catalog's data space is extended, it is to be extended in increments of 5 cylinders.
- VOLUME, which specifies that the catalog is to reside on volume VSER02. Volume VSER02 will be dynamically allocated.
- DATA and INDEX, which specify that VSAM is to allocate 13 cylinders for the catalog itself. VSAM determines the proportion of space that is allocated to the catalog's data and index components. The remainder of the catalog's data space (137 cylinders) is available to contain VSAM objects that reside in a nonunique data space. If the catalog's data component is extended, it is to be extended in increments of 5 cylinders. The catalog's index component is never extended.
- CATALOG, which specifies that the user-catalog connector entry is to be defined in the master catalog, AMASTCAT.CATALOG provides the catalog's update password, MRCATPW2, so that the operator is not prompted to provide it.
- Values and attributes that apply by default to the catalog are:
 - BUFFERSPACE = 3072 bytes
 - ATTEMPTS = 2
 - NOTRECOVERABLE
 - NOWRITECHECK
 - NODESTAGEWAIT
- Null attributes and values of the catalog are:
 - CONTROLPW and READPW passwords
 - CODE
 - AUTHORIZATION
 - OWNER

Define a User Catalog Using the MODEL Parameter: Example 3

In this example, the user catalog defined previously, D27UCAT1, is used as a model for the user catalog being defined, D27UCAT2.

```

//DEFCAT3 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//STEPCAT DD      DSNAME=D27UCAT1,DISP=SHR
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD      *
          DEFINE USERCATALOG( -
              NAME(D27UCAT2) -
              MASTERPW(MASTD27) -
              UPDATEPW(UPDD27) -
              CYLINDERS(50 5) -
              VOLUME(VSER03) -
              MODEL(D27UCAT1/MRPWD27 -
                  D27UCAT1/MRPWD27) ) -
          CATALOG(AMASTCAT/MRCATPW2)
/*

```

The job control language statements are:

- STEPCAT DD, which makes a catalog available for this job step: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE USERCATALOG command defines catalog D27UCAT2. Its parameters are:

- NAME, which names the catalog, D27UCAT2.
- MASTERPW and UPDATEPW, which specify master and update passwords for the catalog: MASTD27 and UPDD27
- CYLINDERS, which specifies that 50 cylinders are to be allocated for the catalog's data space. The catalog itself will reside within this data space. When the catalog's data space is extended, it is to be extended in increments of 5 cylinders.
- VOLUME, which specifies that the catalog is to reside on volume VSER03. Volume VSER03 will be dynamically allocated.
- MODEL, which identifies D27UCAT1 as the catalog to use as a model for D27UCAT2. The attributes and specifications of D27UCAT1 that are not otherwise specified with the above parameters are used to define the attributes and specifications of D27UCAT2. D27UCAT1 contains the catalog entries that describe itself (the master catalog, AMASTCAT, contains a user-catalog connector entry which points to D27UCAT1)—this is why D27UCAT1 is specified as MODEL's catname subparameter. Values and attributes that apply to D27UCAT2 as a result of using D27UCAT1 as a model are:
 - FOR = 365 days (retention period)
 - CYLINDERS = 8 (primary) and 5 (secondary) for the data component and 5 (primary) for the index component
 - BUFFERSPACE = 3072 bytes
 - ATTEMPTS = 2
 - NOTRECOVERABLE
 - NOWRITECHECK
 - NODESTAGEWAIT
 - CONTROLPW and READPW passwords are null
 - CODE is null

- AUTHORIZATION is null
- OWNER is null
- CATALOG, which specifies that the user catalog connector is to be defined in the AMASTCAT catalog. The update password of AMASTCAT is MRCATPW2.

Define a Recoverable User Catalog: Example 4

In this example, a recoverable user catalog is defined.

```
//DEFCAT4 JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//VOL1   DD       VOL=SER=VSER01,UNIT=2314,DISP=OLD
//SYSPRINT DD     SYSOUT=A
//SYSIN  DD       *
          DEFINE USERCATALOG( -
              NAME(D64CAT) -
              MASTERPW(UCATMRPW) -
              UPDATEPW(UCATUPPW) -
              TO(99999) -
              ATTEMPTS(3) -
              CODE(BEQUIET) -
              RECOVERABLE -
              TRACKS(300) -
              FILE(VOL1) -
              VOLUME(VSER01) ) -
          DATA( -
              RECORDS(500 100))
/*
```

The job control language statements are:

- VOL1 DD, which describes the volume on which the catalog is to be defined.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE USERCATALOG command defines a user catalog, D64CAT. Its parameters are:

- NAME, which names the user catalog, D64CAT.
- MASTERPW and UPDATEPW, which specify master and update passwords for the catalog: UCATMRPW and UCATUPPW.
- TO, which specifies that the user catalog is to be retained until it is deleted (that is, the user catalog does not have a retention period).
- ATTEMPTS, which specifies that the operator is prompted for the correct password up to three times.
- CODE, which specifies that the catalog is identified, in the prompting message to the operator, with the name "BEQUIET."
- RECOVERABLE, which specifies that the catalog's volumes contain catalog recovery areas and the LISTCRA, EXPORTRA, IMPORTRA, and RESETCAT commands can be used to restore entries in the catalog.
- TRACKS, which specifies that 300 tracks are to be allocated for the user catalog's data space. No secondary allocation is specified. If additional VSAM space is required on the volume, it can be explicitly defined using the DEFINE SPACE command.

- FILE, which points to the VOL1 DD statement. The VOL1 DD statement directs the OS/VS Scheduler to mount and allocate the volume on which the user catalog is to reside. In this example, the volume is not dynamically allocated.
- VOLUME, which specifies that the user catalog is to reside on volume VSER01.
- DATA, which specifies that VSAM is to allocate space to accommodate 500 records for the user catalog. VSAM determines the total amount of space that is allocated to the catalog's data and index components. The remainder of the catalog's data space is available to contain the catalog recoverable area (CRA) and other VSAM objects that reside in a nonunique data space. If the catalog's data component is extended, it is extended in increments to accommodate 100 records.

You may wish to estimate the amount of space which VSAM will actually allocate to the catalog in order to:

- Ensure that the total amount of space specified via TRACKS is large enough
- Determine how much suballocatable space will be available.

You can make the computation by using the algorithms provided in the section "Estimating the Catalog's Space Requirements." Because the catalog resides on a-2314, the total amount of space will be:

$$aN + b = (.1116 \times 500) + 4 = 59.8 = 60 \text{ tracks}$$

- Values and attributes that apply by default to the catalog are:
 - BUFFERSPACE = 3072 bytes
 - NOWRITECHECK
 - NODESTAGEWAIT
- Null attributes and values of the catalog are:
 - CONTROLPW and READPW passwords
 - AUTHORIZATION
 - OWNER

DELETE

The format of the DELETE command is:

DELETE	<pre>(entryname[/password][b entryname[/password]...]) [ALIAS ALTERNATEINDEX CLUSTER GENERATIONDATAGROUP NONVSAM PAGESPACE PATH SPACE USERCATALOG] [ERASE NOERASE] [FILE(dname)] [FORCE NOFORCE] [PURGE NOPURGE] [SCRATCH NOSCRATCH] [CATALOG(catname[/password])]</pre>
---------------	---

DELETE may be abbreviated: DEL.

DELETE PARAMETERS

Required Parameters

(entryname[/password][b entryname...])
names the entries to be deleted. If more than one entry is to be deleted, the list of entry names must be enclosed in parentheses. This parameter must be the first parameter following DELETE.

entryname
is the name of the entry to be deleted, or the volume serial number of the volume that contains a data space to be deleted. See the section "Generic Name and DELETE" for a description of how to specify a generic name in order to delete multiple entries with one entryname.

If you are deleting a member of a nonVSAM partitioned data set, the entryname must be specified in the format: pdsname(membername).

password
specifies a password for a password-protected entry. Passwords may be specified for each entry name or the catalog's password may be specified through the CATALOG parameter for the catalog that contains the entries to be deleted. Only certain types of catalog entries can be password-protected. To delete a

- Alias, specify the catalog's update- (or higher-level) password.
- Alternate index, specify its master password, or the catalog's master password.
- Cluster, specify its master password, or the catalog's master password.
- Generation Data Group, specify the catalog's update- (or higher-level) password.
- NonVSAM entry, specify the catalog's update- (or higher-level) password.

- Pagespace, specify its master password, or the catalog's master password.
- Path, specify its master password, or the catalog's master password.
- Data space, specify the catalog's update- (or higher-level) password.
- User catalog, specify its master password with the entryname (that is, Access Method Services does not issue a prompting message to the operator or TSO terminal user in this case).

Optional Parameters

**ALIAS|ALTERNATEINDEX|CLUSTER|
GENERATIONDATAGROUP|NONVSAM|PAGESPACE|PATH|
SPACE|USERCATALOG**

specifies the type of the object or entry to be deleted. If the object to be deleted is a catalog or data space, USERCATALOG or SPACE is required.

ALIAS

specifies that the entry to be deleted is an alias entry.

ALTERNATEINDEX

specifies that the object to be deleted is an alternate index and its data and index entries. When a path entry is associated with the alternate index, the path entry is also deleted. When the alternate index has the to-be-upgraded attribute and it is the only such alternate index associated with the base cluster, the base cluster's upgrade-set entry is also deleted.

Abbreviation: AIX

CLUSTER

specifies that the object to be deleted is a cluster, its associated data and index entries, and any related paths and alternate indexes.

Abbreviation: CL

GENERATIONDATAGROUP

specifies that the entry to be deleted is a generation data group entry. A generation data group can be deleted when it is empty, or when you specify the FORCE parameter. **Abbreviation:** GDG

NONVSAM

specifies that the entry to be deleted is a nonVSAM data set entry. If the nonVSAM data set has aliases, all of its alias entries are deleted. If the nonVSAM data set is partitioned, you can delete one of its members by specifying pdsname (member name).

Abbreviation: NVSAM

PAGESPACE

specifies that an inactive page space is to be deleted. A page space is identified as "active" during the operator's IPL procedure.

Abbreviation: PGSPC

PATH

specifies that a path entry is to be deleted. No entries associated with the path are deleted.

SPACE

specifies that all empty VSAM data spaces on a volume are to be deleted. This parameter is required when you want to delete the volume's empty data spaces. A data space can be deleted when it is empty, or when you specify the FORCE parameter. If all data spaces on the volume have been deleted and the volume is not a candidate or catalog volume, its volume entry is also deleted. (Note: When a volume entry in the catalog is deleted, VSAM's ownership of the volume is relinquished.) When the catalog that owns the volume is recoverable; the data space that contains the volume's recovery area is not deleted unless all other data spaces on the volume have been deleted.

Abbreviation: SPC

USERCATALOG

specifies that the object to be deleted is a user catalog. The catalog connector entry in the master catalog is deleted. If the user catalog has aliases, all of the catalog's alias entries in the master catalog are deleted. This parameter must be specified if a user catalog is to be deleted.

A user catalog can be deleted when it is empty (that is, it contains only its self-describing entries and its volume's volume entry), or when you specify the FORCE parameter.

Abbreviation: UCAT

CATALOG(catname[/password])

specifies the name of the catalog that contains the entries to be deleted. See "Order of Catalog Use: DELETE" for information about the order in which catalogs are searched. This parameter cannot be used when a user catalog is to be deleted.

catname

identifies the catalog that contains the entry to be deleted.

password

specifies the password of the catalog that contains the entries to be deleted. If entries to be deleted are password protected and the catalog is also password protected, a password must be supplied either through CATALOG or with the name of each entry to be deleted. The password which must be specified for each entry type is given under the description of the entryname parameter. If a password is supplied through the CATALOG parameter, it takes precedence over any password supplied with the entryname parameter and is used for all entries to be deleted. If a user catalog is to be deleted, a password, if required, must be supplied with the entryname.

Abbreviation: CAT

ERASE|NOERASE

specifies whether the data component of a cluster or alternate index to be deleted is to be erased, that is, overwritten with binary zeros. This parameter will override whatever was coded when the cluster or alternate index was defined or last altered. This parameter should be specified only when a cluster or alternate index entry is to be deleted. When you specify ERASE, you must identify the to-be-deleted entry's catalog with a JOBCAT or STEPCAT DD statement unless

- The entry is in the master catalog, or

- The first qualifier in the entry's qualified name identifies the catalog (that is, the first qualifier is the catalog's name or alias).

ERASE

specifies that the data component is to be overwritten with binary zeros when the cluster or alternate index is deleted. If ERASE is specified, the volume that contains the data component must be mounted.

Abbreviation: ERAS

NOERASE

specifies that the data component is not to be overwritten with binary zeros when the cluster or alternate index is deleted.

Abbreviation: NERAS

FILE(dname)

specifies the name of the DD statement that identifies:

- The volume that contains a unique data set to be deleted.
- The volume that contains a nonVSAM data set to be deleted (with SCRATCH).
- The entry (data set) to be deleted if erase processing is to be performed. (**Note:** For erase processing, a JOBCAT or STEPCAT DD statement is required unless the data set is cataloged in the master catalog or its high level qualifier is the name or alias of the catalog containing the entry.)
- The volume that contains the data space to be deleted.
- The partitioned data set from which a member (or members) is to be deleted.
- The catalog recovery volume(s) when the entry being deleted is in a recoverable catalog. If the volumes are of a different device type, concatenated DD statements must be used. (The catalog recovery volume is the volume whose recovery space contains a copy of the entry being deleted. See "Deleting Catalog Entries" for more details on catalog recovery volume mounting requirements.)
- The volume(s) owned by the catalog when you delete a catalog using the FORCE option.

When more than one volume is to be identified (for example, a multivolume data set), FILE identifies the DD statement which specifies all volumes. If, in any of the above cases the volumes are of a different device type, concatenated DD second DELETE statements must be used.

When deleting multivolume nonVSAM data sets with the SCRATCH option, DELETE SCRATCH processing will require access to each volume in the entry's catalog record before the scratch can be issued. This will require either all volumes to be mounted, online, and allocatable to the job, or the use of the FILE parameter specifying a DD card allocating at least one mountable unit (not permanently resident or reserved). Deferred mount must be specified on the DD card so that allocation will flag the UCB to allow demount/mount requests to be issued for the unit as required during delete processing. If access to the volumes cannot be provided, then DELETE NOSCRATCH can be used to uncatalog the nonVSAM data set and the user will assume the responsibility of scratching the format-1 DSCBs from all the volumes.

In the following example nonVSAM data set 'TWO.TEST' exists on volumes TEST01, TEST02, TEST03, TEST04, TEST05, and TEST06. Only the first three volumes are mounted and online but TEST03 resides on a mountable unit.

```
//DLTB      EXEC  PGM=IDCAMS, REGION=400K
//STEP CAT  DD    DSN=USERCAT, DISP=SHR
//SYS PRINT DD    SYSOUT=A
//DD1       DD    DSN=TWO.TEST, DISP=OLD,
//           UNIT=(3330,3,DEFER),
//           VOL=SER=(TEST01,TEST02,TEST03)
//SYSIN     DD    *
          DELETE -
            TWO.TEST -
            NONVSAM -
            FILE(DD1) -
            PURGE -
            CATALOG(USERCAT)
/*
```

If you do not specify FILE and VSAM requires access to a volume or volumes during the delete processing, the volume(s) will be dynamically allocated. When erase processing is to be performed, the entryname will be dynamically allocated if FILE is not specified. Dynamic allocation requires that the volume(s) must be mounted as permanently resident or reserved.

FILE is required when deleting a nonempty user catalog using the FORCE option, and the catalog has data spaces residing on volumes other than the volume on which the catalog resides.

FORCE|NOFORCE

The FORCE parameter allows you to delete data spaces, generation data groups, and VSAM user catalogs without first ensuring that these objects are empty.

- When you delete space using the FORCE parameter, all VSAM data spaces on the volume(s) are scratched from the VTOC. VSAM's ownership of the volume(s) is given up. Before the data spaces are scratched, VSAM verifies that no VSAM clusters on the volume are open. The clusters remain cataloged, but their entries are marked unusable. Subsequently, you can delete the cluster from the catalog.

You cannot specify FORCE to delete the space on a volume that contains the VSAM master catalog or a user catalog. When you specify FORCE, you must supply the master password of the catalog that contains the volume's entry.

- When you delete a user catalog using the FORCE parameter, the catalog is deleted even though it contains catalog entries for objects that haven't been deleted. All data spaces are deleted from each volume owned by the catalog. Before the data spaces are scratched, VSAM verifies that no VSAM clusters on the volume are open. The volumes are available to be used by the OS/VS2 MVS system or to be assigned to other VSAM catalogs. All cataloged objects are automatically deleted, but the contents of each cluster and alternate index are not erased. NonVSAM data set entries in the catalog are deleted, but the nonVSAM data set's space on the volume is undisturbed. The nonVSAM data set can be located with its DSCB in the volume's VTOC. VSAM clusters and alternate indexes, even though they might reside in unique data spaces, cannot be located, because the data space's DSCB is scratched when the catalog is deleted.

When you delete a user catalog using the FORCE parameter, you must supply the master password of the catalog being deleted.

- When you delete a generation data group using FORCE, the GDG entry is deleted even though it might point to generation data sets (that is, it might point to nonVSAM entries in the catalog). Each nonVSAM data set entry (that is, the entry for a generation data set) pointed to by the GDG base entry is deleted before the GDG base entry is deleted. However, the nonVSAM data set's space and contents on the volume is undisturbed. The data set can be located with its DSCB in the volume's VTOC.

NOFORCE, the default option, causes the DELETE command to terminate when you request the deletion of a nonempty data space, generation data group, or catalog.

Abbreviations: FRC and NFRC

PURGE|NOPURGE

specifies whether the entry is to be deleted regardless of the retention period specified. This parameter cannot be used if a data-space entry is to be deleted.

PURGE

specifies that the entry is to be deleted even if the retention period, specified in the TO or FOR parameter, has not expired.

Abbreviation: PRG

NOPURGE

specifies that the entry is not to be deleted if the retention period has not expired. If NOPURGE is coded and the retention period has not expired, the entry is not deleted.

Abbreviation: NPRG

SCRATCH|NOSCRATCH

has two uses:

1. With NONVSAM, SCRATCH|NOSCRATCH specifies whether a nonVSAM data set is to be scratched—removed from the VTOC—of the volume on which it resides.
2. With CLUSTER, ALTERNATEINDEX, SPACE, or PAGESPACE, NOSCRATCH specifies that the entry is to be deleted from the catalog without access to the volume that contains the object defined by the entry. NOSCRATCH cannot be specified for VSAM clusters, alternate indexes, or page spaces in recoverable catalogs. See the section "VSAM Catalog Cleanup" in the chapter "Data Security and Protection" for important information about the use of DELETE NOSCRATCH.

NOSCRATCH will be ignored when deleting the VSAM cluster entry (created for the source catalog) on the receiving catalog's volume as a result of the copy catalog operation. See the section "Copy-Catalog Procedure" in the chapter "Copying and Printing" for more detailed information.

When you specify SCRATCH, you must identify the to-be-deleted entry's catalog with a JOBCAT or STEPCAT DD statement, unless:

- The entry is in the master catalog, or
- The first qualifier in the entry's qualified name identifies the catalog (that is, the first qualifier is the catalog's name or alias).

You should specify NOSCRATCH if the format-1 DSCB of a nonVSAM data set has already been scratched from the VTOC.

You should specify NOSCRATCH if you are deleting a nonVSAM data set which was defined with an esoteric device type (for example, SYSDA). If you specify SCRATCH, your DELETE will fail. You should use the SCRATCH function of the OS/VS utility, IEHPRGM, to remove the format-1 DSCB from the VTOC.

Abbreviations: SCR and NSCR

DELETE EXAMPLES

Deleting a Key-Sequenced VSAM Cluster: Example 1

In this example, a key-sequenced cluster is deleted. Alternate indexes and paths related to the key-sequenced cluster are deleted automatically by Access Method Services. Access Method Services will dynamically allocate the key-sequenced data set so that the data can be overwritten (as specified by the ERASE option).

```
//DELET1  JOB    ...
//STEP1   EXEC  PGM=IDCAMS
//JOB CAT DD   DSN=D27UCAT2,DISP=OLD
//SYS PRINT DD  SYSOUT=A
//SYS IN  DD   *
          DELETE -
            D40.EXAMPLE.KSDS1 -
            PURGE -
            ERASE -
            CATALOG(D27UCAT2/MRPWD27)
/*
```

The job control language statement is:

- JOBCAT DD, which makes a catalog available for the job: D27UCAT2.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command deletes the key-sequenced VSAM cluster from the D27UCAT2 catalog. Its parameters are:

- D40.EXAMPLE.KSDS1, which is the entryname of the object being deleted. D40.EXAMPLE.KSDS1 is a key-sequenced VSAM cluster.
- PURGE, which specifies that the cluster is to be deleted regardless of its retention period or date.
- ERASE, which specifies that the cluster's data component is to be overwritten with binary zeros, regardless of a NOERASE attribute that might have been specified when the cluster was defined or altered.
- CATALOG, which identifies the catalog that contains the cluster's entries, D27UCAT2. The catalog's master password is required, unless the entry's master password is specified with the entryname.

Deleting an Entry-Sequenced VSAM Cluster and All Empty Data Spaces On a Volume: Example 2

In this example, an entry-sequenced VSAM cluster is deleted. Alternate indexes and paths related to the entry-sequenced cluster are automatically deleted by Access Method Services. Next, all empty VSAM data spaces on volume VSER03 are deleted and the volume's volume entry is deleted.

```
//DELET2 JOB ...
//JOB CAT DD DSNAME=D27UCAT2,DISP=OLD
//STEP1 EXEC PGM=IDCAMS
//DD2 DD VOL=SER=VSER03,UNIT=3330,DISP=OLD
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
DELETE -
EXAMPLE.ESDS2/DEPT26M -
PURGE -
CLUSTER
DELETE -
VSER03 -
SPACE -
FILE (DD2)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for the job: D27UCAT2.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The first DELETE command deletes the only VSAM cluster on volume VSER03. The DELETE command's parameters are:

- EXAMPLE.ESDS2, which is the entryname of the cluster to be deleted, and DEPT26M, which is the cluster's master password.
- PURGE, which specifies that the cluster is to be deleted regardless of its retention period or date. Alternate indexes and paths related to the entry-sequenced cluster are also deleted regardless of their retention periods or dates.
- CLUSTER, which specifies that the entryname EXAMPLE.ESDS2 identifies a VSAM cluster (that is, the entryname identifies a cluster entry).

The second DELETE command causes VSAM to examine each data space cataloged in the volume's entry and, if the data space is empty, delete the data space. VSAM deletes the volume entry when there are no data spaces cataloged in it and when there are no data sets identified that specify the volume as a candidate volume. VSAM also updates the volume's VTOC to reflect the deleted data spaces. The DD2 DD statement ensures that volume VSER03 is mounted when its data spaces are deleted. In this example JCL has been used instead of allowing Access Method Services to dynamically allocate the volume. If the volume were mounted as permanently RESIDENT or RESERVED, Access Method Services could dynamically allocate the volume without the need of JCL or the FILE parameter. The DELETE command's parameters are:

- VSER03, which identifies the volume that contains empty VSAM data spaces.
- SPACE, which specifies that the entryname identifies a volume entry. When a volume's data spaces are to be deleted, the SPACE parameter is required.
- FILE, which specifies the dname of a DD statement that describes the VSAM volume and causes it to be mounted.

Deleting Two Key-Sequenced Clusters: Example 3

In this example, two key-sequenced clusters, D40.MYDATA and ENTRY, are deleted from a recoverable catalog. This example illustrates how more than one cataloged object is deleted with a single DELETE command.

```
//DELET3  JOB    ...
//JOB CAT DD    DSNAME=MYCAT,DISP=SHR
//STEP1   EXEC  PGM=IDCAMS
//SYS PRINT DD  SYSOUT=A
//SYS IN  DD    *
        DELETE -
            (D40.MYDATA -
             ENTRY) -
            PURGE -
            CLUSTER
/*
```

The job control language statements are:

- JOBCAT DD, which makes the catalog MYCAT available for the job.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command deletes the key-sequenced clusters D40.MYDATA and ENTRY. The high-level qualifier of the object D40.MYDATA causes the catalog D27UCAT1 to be searched, since it is assumed that D40 is an alias of the catalog D27UCAT1. The object ENTRY is found in the catalog MYCAT, which is made available to the job through the JOBCAT DD statement. Both entries will be dynamically allocated so that their respective catalog recovery areas can be updated. The DELETE command's parameters are:

- D40.MYDATA and ENTRY, which identify the objects to be deleted. D40.MYDATA and ENTRY are the entrynames of two key-sequenced clusters. It is assumed that neither cluster is password-protected at this time. If either cluster is password-protected, the operator is prompted to supply the cluster's master password.
- PURGE, which specifies that the cluster is to be deleted regardless of its retention period or date.
- CLUSTER, which specifies that D40.MYDATA and ENTRY are clusters (that is, the names identify cluster catalog records).

Deleting a NonVSAM Data Set's Entry: Example 4

In this example, a nonVSAM data set's entry (cataloged in a VSAM user catalog) is deleted. The SCRATCH parameter is implied (that is, it is the default). A FILE parameter and its associated DD statement are provided to allocate the data set's volume. In this example, dynamic allocation is not used to provide catalog or volume allocation.

```

//DELET4 JOB ...
//JOB CAT DD DSN= D27UCAT1,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//DD1 DD VOL=SER=VSER02,UNIT=2314,DISP=OLD,
// DSN=EXAMPLE.NONVSAM
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
DELETE -
EXAMPLE.NONVSAM -
FILE (DD1) -
PURGE -
CATALOG(D27UCAT1/UPPWD27)
/*

```

The job control language statements are:

- JOB CAT DD, which makes a catalog available for this job: D27UCAT1.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command deletes the nonVSAM data set EXAMPLE.NONVSAM. The DELETE command's parameters are:

- EXAMPLE.NONVSAM, which is the entryname of the object to be deleted.
- FILE, which specifies the dname of a DD statement that describes the nonVSAM data set's volume and causes it to be mounted. When the data set is deleted, its DSCB entry in the volume's VTOC is removed.
- PURGE, which specifies that the nonVSAM data set's retention period or date is to be ignored. If PURGE is not specified and the nonVSAM data set's retention period has not yet expired, VSAM wouldn't delete its entry.
- CATALOG, which identifies the catalog that contains the entries, D27UCAT1, and its update password, UPPWD27.

Deleting Nonempty VSAM Data Spaces on a Volume: Example 5

In this example, all of the data spaces on volume VSER05 are deleted, whether or not they are empty. All VSAM data spaces on the volume are scratched from the VTOC and the ownership of the volume is surrendered by the catalog, AMASTCAT. Any clusters or alternate indexes contained within the data spaces on volume VSER05 are marked unusable in the catalog. Because all VSAM data spaces on the volume are deleted, the volume's catalog entry is also deleted.

```

//DELET5 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
DELETE -
VSER05 -
SPACE -
FORCE -
CATALOG(AMASTCAT/MRCATPW2)
/*

```

The job control language statements are:

- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command's parameters are:

- VSERO5, which identifies the volume with its serial number. Because the volume's VTOC must be updated, Access Method Services will dynamically allocate the volume.
- SPACE, which specifies that a volume entry is to be modified or deleted. When a volume's data space is to be deleted, the SPACE parameter is required.
- FORCE, which specifies that all VSAM data spaces on the volume are to be deleted, whether or not they are empty, and that volume ownership is to be released.
- CATALOG, which identifies the catalog that owns the volume, AMASTCAT. The catalog's update (or higher level) password is required.

Deleting a User Catalog: Example 6

In this example, a user catalog is deleted. A user catalog can be deleted when it is empty—that is, when there are no objects except the catalog's volume cataloged in it. If the catalog is not empty, it cannot be deleted unless the FORCE parameter is specified.

```
//DELET6 JOB      ...
//STEP1  EXEC    PGM=IDCAMS
//SYSPRINT DD    SYSOUT=A
//SYSIN   DD     *
        DELETE -
            D27UCAT1/MRPWD27 -
            PURGE -
            USERCATALOG
/*
```

The DELETE command deletes the catalog. The volume that contained the catalog is now available for ownership by another VSAM catalog. The DELETE command also deletes the catalog's user catalog connector entry in the master catalog. The DELETE command's parameters are:

- D27UCAT1, the name of the user catalog. The catalog's update (or higher level) password is required.
- PURGE, which specifies that the user catalog's retention period or date is to be ignored. If PURGE is not specified and the catalog's retention period has not yet expired, it won't be deleted.
- USERCATALOG, which identifies D27UCAT1 as a user catalog.

Deleting an Alias Entry: Example 7

In this example, an alias entry, EXAMPLE.NONVSAM1, is removed from catalog D27UCAT1.

```
//DELET7 JOB      ...
//JOB CAT DD     DSNAME=D27UCAT1,DISP=OLD
//STEP1  EXEC    PGM=IDCAMS
//SYSPRINT DD    SYSOUT=A
//SYSIN   DD     *
        DELETE -
            EXAMPLE.NONVSAM1 -
            ALIAS -
            CATALOG(D27UCAT1/MRPWD27)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command removes an alias entry from catalog D27UCAT1. Its parameters are:

- EXAMPLE.NONVSAM1, the entryname of the object to be deleted. EXAMPLE.NONVSAM1 identifies an alias entry.
- ALIAS, which specifies the type of entry to be deleted. VSAM verifies that EXAMPLE.NONVSAM1 is an alias entry, then deletes it. If EXAMPLE.NONVSAM1 incorrectly identified another entry by mistake, VSAM would not delete the entry, but would note the discrepancy with a message to the programmer.
- CATALOG, which identifies the catalog that contains the entry, D27UCAT1, and its master password, MRPWD27.

Deleting Generically-Named Entries: Example 8

In this example, each catalog entry with the name "GENERIC.*.ABLE" is deleted, where "*" is any 1- to 8-character simple name. The name "GENERIC.*.ABLE" is a generic name, and this example shows how all catalog entries with the same generic name are deleted.

```
//DELET8 JOB ...
//JOBCAT DD DSNAME=D27UCAT1,DISP=OLD
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DELETE -
        GENERIC.*.ABLE -
PURGE -
        CATALOG(D27UCAT1/MRPWD27)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command removes all entries (and their associated entries) with the generic name "GENERIC.*.ABLE" from catalog D27UCAT1. Its parameters are:

- GENERIC.*.ABLE, a generic name, identifies all catalog entries with the name GENERIC.*.ABLE.
- PURGE, which specifies that each entry is to be purged regardless of the retention period or date specified when it was defined.
- CATALOG, which identifies the catalog that contains the entries, D27UCAT1, and its master password, MRPWD27.

List a Generation Data Group's Entries, Then Delete the Group and Its Data Sets: Example 9

In this example, a generation data group, GDG01, and its associated (generation data set) entries are listed. Next, the only generation data set in the group is deleted. Finally, the generation data group base catalog entry is deleted.

```
//DELET9   JOB    ...
//STEP1   EXEC   PGM=IDCAMS
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD     *
LISTCAT -
        ENTRIES(GDG01) -
        ALL
DELETE -
        GDG01.G0001V00 -
        PURGE
DELETE -
        GDG01 -
        GENERATIONDATAGROUP -
        PURGE
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

Because no catalog was specified in a CATALOG parameter or with a JOBCAT or STEPCAT DD statement, VSAM assumes all entries required for this command sequence are in the master catalog.

The LISTCAT command lists the generation data group GDG01 and its associated generation data set entries. Its parameters are:

- ENTRIES, which specifies that the entry GDG01 is to be listed. Since the entry GDG01 is a generation data group entry, its associated generation data set's (nonVSAM) entries are also listed. In addition, if one of the generation data sets has aliases, the alias entries associated with the generation data set's entry are listed.
- ALL, which specifies that all fields are to be listed.

The first DELETE command removes the nonVSAM data set entry for the only generation data set in the generation data group: GDG01.G0001V00. Its parameters are:

- GDG01.G0001V00, the entryname of the object being deleted. GDG01.G0001V00 identifies the only generation data set in the generation data group GDG01.
- PURGE, which specifies that the generation data set's retention period or date is to be ignored. If PURGE is not specified and the generation data set's retention period has not yet expired, VSAM wouldn't delete its entry.

The second DELETE command removes the generation data group base catalog entry from the catalog. Its parameters are:

- GDG01, the entryname of the object being deleted. GDG01 identifies the generation data group base entry.
- GENERATIONDATAGROUP, which specifies the type of entry to be deleted. VSAM verifies that GDG01 is a generation data group entry, then deletes it. If GDG01 incorrectly specified another type of entry, VSAM would not delete the entry, but would note the discrepancy with a message to the programmer.

- PURGE, which specifies that the generation data group's retention period or date is to be ignored. If PURGE is not specified and the generation data group's retention period has not yet expired, VSAM wouldn't delete its entry.

Deleting a Member of a Partitioned (NonVSAM) Data Set: Example 10

In this example, the MEM1 member of partitioned data set EXAMPLE.NONVSAM2 is deleted. Next, the nonVSAM data set itself is deleted.

```
//DELET10   JOB      ...
//JOB CAT   DD      DSN= D27UCAT1,DISP=SHR
//STEP1     EXEC    PGM=IDCAMS
//SYS PRINT DD      SYSOUT=A
//SYS IN    DD      *
DELETE -
          EXAMPLE.NONVSAM2(MEM1) -
          CATALOG(D27UCAT1/MRPWD27)
DELETE -
          EXAMPLE.NONVSAM2 -
          PURGE -
          CATALOG(D27UCAT1/MRPWD27)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The first DELETE command deletes one of the members of a partitioned data set, EXAMPLE.NONVSAM2(MEM1), from the user catalog D27UCAT1. Its parameters are:

- EXAMPLE.NONVSAM2(MEM1), which is the entryname of one of the members of the partitioned (nonVSAM) data set, EXAMPLE.NONVSAM2. The entryname identifies the object to be deleted.
- CATALOG, which identifies the catalog that contains the entry, D27UCAT1, and its master password, MRPWD27.

The second DELETE command deletes all remaining members, as well as the data set itself, of the partitioned nonVSAM data set EXAMPLE.NONVSAM2. Its parameters are:

- EXAMPLE.NONVSAM2, which is the entryname of the object to be deleted.
- PURGE, which specifies that the nonVSAM data set's retention period or date is to be ignored. If PURGE had not been specified and the nonVSAM data set's retention period has not yet expired, VSAM wouldn't delete its entry.
- CATALOG, which identifies the catalog that contains the entry, D27UCAT1, and its master password, MRPWD27.

In the second part of this example, the DSCB entry in the volume's VTOC is removed. Dynamic allocation is used to allocate the data set's volume.

Deleting a Page Space: Example 11

In this example, page space SYS1.PAGE2 is deleted.

```
//DELET11    JOB    ...
//STEP1     EXEC   PGM=IDCAMS
//SYSPRINT  DD     SYSOUT=A
//SYSIN     DD     *
           DELETE -
             SYS1.PAGE2 -
             PURGE -
             PAGESPACE
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DELETE command removes a page space entry, SYS1.PAGE2, from the master catalog. Its parameters are:

- SYS1.PAGE2, the entryname of the object to be deleted. SYS1.PAGE2 identifies a page space entry.
- PURGE, which specifies that the page space entry is to be deleted regardless of the retention period or date specified when it was defined.
- PAGESPACE, which specifies the type of entry to be deleted. VSAM verifies that SYS1.PAGE2 is a page space entry, then deletes it. If SYS1.PAGE2 incorrectly identified another type of entry, VSAM would not delete it, but would send an error message to the programmer.

EXPORT

The format of the EXPORT command, when it is used to disconnect a user catalog, is:

EXPORT	usercatname[/password] DISCONNECT
--------	--------------------------------------

The format of the EXPORT command, when it is used to move a cluster or alternate index, is:

EXPORT	entryname[/password] {OUTFILE(dname) OUTDATASET(entryname)} [INFILE(dname)] [ERASE NOERASE] [INHIBITSOURCE NOINHIBITSOURCE] [INHIBITTARGET NOINHIBITTARGET] [PURGE NOPURGE] [TEMPORARY PERMANENT]
--------	--

EXPORT PARAMETERS

Required Parameters

entryname[/password]

is a required parameter that names the cluster, alternate index, or user catalog to be exported. This parameter must be the first parameter following EXPORT. If you are exporting a user catalog, you must supply the update (or higher level) password of the master catalog. If you are exporting a cluster or alternate index, you must supply its (the object's) master password. If any password-protected paths are defined over the cluster or alternate index being exported, you must supply the master password of the catalog containing the objects in order to export the protection attributes of all the paths.

Abbreviation: EXP

DISCONNECT

specifies that a user catalog is to be exported. The entry for the user catalog will be deleted from the master catalog. Also, the user catalog's alias entries are deleted from the master catalog.

A JOBCAT/STEP CAT can be used to delete a user catalog entry from a user catalog. Any alias entries will also be deleted.

DISCONNECT is a required parameter if EXPORT is issued for a user catalog. The volume that contains the user catalog can be physically moved to the system to which the catalog will be imported.

If DISCONNECT is coded, the entry name and the master catalog's update or higher-level password, and no other parameters, must be coded with it. To make a user catalog available in other systems and in the original system, code the IMPORT command to import the user catalog to each system to which it is to be available, but do not EXPORT the user catalog.

Abbreviation: DCON

OUTFILE(dname)

specifies the name of the DD statement that identifies the output data set to be created as a result of the EXPORT

command. With the exception of blocksize, the data set characteristics of the output data set that is to contain the cluster to be exported should not be specified. The characteristics of the output data set are described in the section "Moving Entries."

Blocksize is defaulted by the EXPORT command to 2048. Blocksize may be specified in the DD statement to override this default and improve performance.

Abbreviation: OFILE

OUTDATASET(entryname)
specifies the name of the data set that is to receive the data being exported. If OUTDATASET is specified, the entryname is dynamically allocated. The characteristics of the output data set are described in the section "Moving Entries."

Abbreviation: ODS

Optional Parameters

ERASE|NOERASE
specifies whether the data component of the cluster or alternate index to be exported is to be erased (that is, overwritten with binary zeros). This parameter overrides whatever was specified when the object was defined or last altered. This parameter can be specified only if the object is to be permanently exported (that is, deleted from the original system).

Abbreviations: ERAS and NERAS

INFILE(dname)
specifies the name of the DD statement that identifies the cluster or alternate index to be exported. If the cluster or alternate index has been defined with a maximum logical record length greater than 32,760 bytes, EXPORT processing terminates with an error message.

In addition to the DD statement for INFILE, you must identify the entry's catalog with a JOBCAT or STEPCAT DD statement unless:

- The object's entry is in the master catalog, or
- The first qualifier of the object's name identifies a catalog (that is, the first qualifier is the catalog's name or alias).

When INFILE and its DD statement aren't specified for a to-be-exported object, it is dynamically allocated.

Abbreviation: IFILE

INHIBITSOURCE|NOINHIBITSOURCE
specifies whether the original data records (that is, the data records of the source cluster or alternate index) can be accessed for any operation other than retrieval after its copy is exported. This specification can later be altered through the ALTER command.

INHIBITSOURCE
specifies that the source object in the original system cannot be accessed for any operation other than retrieval. This parameter can be specified only when the object is to be temporarily exported (that is, a backup copy of the object is made; the object itself remains in the original system).

Abbreviation: INHS

NOINHIBITSOURCE

specifies that the source object in the original system can be accessed for any kind of operation.

Abbreviation: NINHS

INHIBITTARGET|NOINHIBITTARGET

specifies whether the copy's records (that is, the target alternate index or cluster) can be accessed for any operation other than retrieval after it has been imported to another system. This specification can be altered through the ALTER command.

INHIBITTARGET

specifies that the target object cannot be accessed for any operation other than retrieval after it has been imported into another system.

Abbreviation: INHT

NOINHIBITTARGET

specifies that the target object can be accessed for any type of operation after it has been imported into another system.

Abbreviation: NINHT

PURGE|NOPURGE

specifies whether the cluster or alternate index to be exported is to be deleted from the original system regardless of the retention period, specified in a TO or FOR parameter when the object was defined. This parameter can be specified only if the object is to be permanently exported, that is, deleted from the original system.

PURGE

specifies that the object is to be deleted even if the retention period has not expired.

Abbreviation: PRG

NOPURGE

specifies that the object is not to be deleted unless the retention period has expired.

Abbreviation: NPRG

TEMPORARY|PERMANENT

specifies whether the cluster or alternate index to be exported is to be deleted from the original system.

TEMPORARY

specifies that the cluster or alternate index is not to be deleted from the original system. The object in the original system is marked as temporary to indicate that another copy exists and that the original copy can be replaced. To replace the original copy, a portable copy created by an EXPORT command must be imported to the original system. The IMPORTJ command deletes the original copy, defines the new object, and copies the data from the portable copy into the newly defined object.

Abbreviation: TEMP

PERMANENT

specifies that the cluster or alternate index is to be deleted from the original system. Its storage space is freed. If its retention period has not yet expired, you must also code PURGE.

Abbreviation: PERM

EXPORT EXAMPLES

Exporting a Key-Sequenced Cluster: Example 1

In this example, a key-sequenced cluster, D40.EXAMPLE.KSDS1, is exported from a user catalog, D27UCAT1. The cluster is copied to a portable file, TAPE2, and its catalog entries are modified to prevent the cluster's data records from being updated, added to, or erased.

```
//EXPORT1 JOB    ...
//STEP1  EXEC   PGM=IDCAMS
//RECEIVE DD    DSNAME=TAPE2,UNIT=(2400-3,,DEFER),
//          DISP=NEW,VOL=SER=003030,
//          DCB=(BLKSIZE=6000,DEN=3),LABEL=(1,SL)
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          EXPORT -
            D40.EXAMPLE.KSDS1 -
            OUTFILE(RECEIVE) -
            TEMPORARY -
            INHIBITSOURCE
/*
```

The job control language statements are:

- RECEIVE DD, which describes the portable file, a magnetic tape file, that is to receive a copy of the cluster's records. The DCB BLKSIZE parameter overrides the EXPORT default of 2048 to improve performance.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The EXPORT command copies the key-sequenced cluster, D40.EXAMPLE.KSDS1, to a portable file, TAPE2, and modifies the cluster's catalog entries (its cluster, data, and index entries). Access Method Services dynamically allocates D40.EXAMPLE.KSDS1. Since the high-level qualifier of the cluster name is the name of the alias of D27UCAT2, the cluster can be allocated without specifying a JOBCAT or STEPCAT DD statement. The EXPORT command's parameters are:

- D40.EXAMPLE.KSDS1, which identifies the cluster to be exported. Because the cluster is not password protected, no password is provided with the cluster's entryname.
- OUTFILE, which points to the RECEIVE DD statement. The RECEIVE DD statement describes the portable file, TAPE2, that is to contain a copy of the cluster.
- TEMPORARY, which specifies that the cluster is not to be deleted. The cluster's catalog entry is marked "temporary" to indicate that another copy of the cluster exists and that the original copy can be replaced. (See the IMPORT Example 2, "Importing a Key-Sequenced Cluster.")
- INHIBITSOURCE, which specifies the cluster (that is, the copy of it that remains in the original system, as a result of TEMPORARY) cannot be modified. User programs are allowed only to read the cluster's records.

Exporting an Entry-Sequenced Cluster: Example 2

In this example, an entry-sequenced cluster is exported to a portable file and then it is deleted from the system.

```

//EXPORT2 JOB ...
//STEP1 EXEC PGM=IDCAMS
//RECEIVE DD DSNAME=TAPE1,UNIT=(2400-3,,DEFER),
// VOL=SER=001147,LABEL=(1,SL),DISP=NEW
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
EXPORT -
        D50.EXAMPLE.ESDS1/DEPT26M -
        OUTFILE(RECEIVE) -
        PURGE
/*

```

The job control language statements are:

- RECEIVE DD, which describes the portable file, TAPE1, that is to contain a copy of the exported entry-sequenced cluster.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The EXPORT command copies the entry-sequenced cluster, D50.EXAMPLE.ESDS1, to a portable file, TAPE1. The cluster is deleted from the system after it is copied into the portable file. Access Method Services dynamically allocates D50.EXAMPLE.ESDS1. Since the high-level qualifier of the cluster is the name of the alias of the catalog D27UCAT2, the cluster can be dynamically allocated without specifying a JOBCAT or STEPCAT DD statement. The EXPORT command's parameters are:

- D50.EXAMPLE.ESDS1, which identifies the entry-sequenced cluster to be exported, and which specifies the cluster's master password, DEPT26M. When a cluster is exported, its master password is required if it is password protected.
- OUTFILE, which points to the RECEIVE DD statement. The RECEIVE DD statement describes the portable data set, TAPE1, that is to receive a copy of the cluster.
- PURGE, which allows the cluster to be deleted regardless of its retention period or date.

Because TEMPORARY is not specified, Access Method Services assumes the cluster is to be deleted once TAPE1 contains a copy of it.

Because INHIBITTARGET is not specified, Access Method Services assumes the cluster can be updated (by users of the other system) when it is imported to another system.

Exporting a User Catalog: Example 3

In this example, the user catalog D27UCAT1 is exported—that is, it is disconnected from the system. Its cataloged objects are no longer available to users of the system.

```

//EXPORT3 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
EXPORT -
        D27UCAT1/MRPWD27 -
        DISCONNECT
/*

```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The EXPORT command removes the user catalog connector entry for D27UCAT1 from the master catalog. The catalog becomes unavailable to system users until the system programmer reconnects it to the system, using an IMPORT CONNECT command. The EXPORT command's parameters are:

- D27UCAT1, which identifies the object to be exported. When a user catalog is exported, the master catalog's master password is required. (See IMPORT Example 1, "Import a User Catalog.")
- DISCONNECT, which identifies the exported object as a user catalog. When a user catalog is exported, DISCONNECT is required.

EXPORTRA

The format of the EXPORTRA command is:

EXPORTRA	OUTFILE(dname) CRA((dname1 {ALL[b INFILE(dname2)]} ENTRIES((entryname[b dname3]) [(entryname[b dname3])...]] NONE) [(dname1...)] [FORCE NOFORCE] [MASTERPW(password)]
-----------------	---

EXPORTRA may be abbreviated: XPRA.

EXPORTRA PARAMETERS

Required Parameters

OUTFILE(dname)
identifies the DD statement that describes the data set (usually a movable disk pack or magnetic tape reel) that is to contain a copy of the VSAM objects and catalog entries identified with the CRA parameter. With the exception of blocksize, the data set characteristics of the output data set that is to contain the cluster to be exported should not be specified. The characteristics of the output data set are described in the section "Copying a Catalog Entry from the Catalog Recovery Area." Blocksize is defaulted by the EXPORTRA command to 2048. Blocksize may be specified in the DD statement to override this default and improve performance.

Abbreviation: OFILE

CRA(
 (dname1
 {ALL[b INFILE(dname2)]}
 ENTRIES(
 (entryname[b dname3])
 [(entryname[b dname3])...]]
 NONE)
 [(dname1...)]

The CRA parameter and its subparameters describe the catalog recovery area to be referred to, identify the specific entries to be copied from the catalog recovery area, and specify the volumes that contain the entries' data. The subparameters are repeated to describe each catalog recovery area to be referred to, and whether its entries are to be copied.

dname1
identifies the DD statement that describes a volume whose catalog recovery area is to be referred to. Each dname1 you specify is followed by one of the following subparameters which specifies the action to be taken with regard to the catalog recovery area:

ALL

specifies that each entry in dname1's catalog recovery area is to be copied to the OUTFILE data set. In addition, the contents (that is, data records in logical sequential order) of each VSAM cluster and alternate index are copied, following the copy of the catalog recovery area entry. If any of these VSAM clusters or alternate indexes

has data which resides on volumes other than that identified via dname1, you use INFILE to describe all the volumes.

INFILE(dname2)
 identifies the DD statement that describes multiple volumes to be used in the recovery process. INFILE is required if the VSAM clusters and/or alternate indexes whose entries are contained in dname1's catalog recovery area reside on multiple volumes, in addition to the dname2 DD statement, including the volume referenced by the dname1 DD statement.

Abbreviation: IFILE

ENTRIES(
 (entryname[b dname3])
 [(entryname[b dname3](...)]
 names each entry in the catalog recovery area that is to be copied to the OUTFILE data set. All other entries in the catalog recovery area are not copied.

entryname
 names an object whose catalog entry is contained in the dname1 volume's catalog recovery area. You can specify the entryname of the following entry types: alternate index, cluster, and nonVSAM.

When you specify a cluster entry, the cluster's data and index entries and all associated path entries are also copied. In addition, the cluster's contents (that is, its data records, in logical sequential order) are copied, following the copies of the cluster's catalog entries.

When you specify an alternate index entry, the alternate index's entries and contents and all associated path entries are copied as though it were a key-sequenced cluster.

dname3
 identifies the DD statement that describes each volume that contains part of the cluster or alternate index. dname3 is ignored when specified for a nonVSAM entry.

dname3 is required when the entry is a VSAM cluster or alternate index whose data resides on volumes other than the volume identified with the dname1 DD statement. The dname3 DD statement must also include the volume referenced by the dname1 DD statement.

Abbreviation: ENT

NONE
 specifies that no object in the volume's catalog recovery area is to be recovered. You must specify dname1 NONE for each volume referenced in a dname2 or dname3 DD statement if you do not wish the entries in its catalog recovery area to be recovered.

Optional Parameters

FORCE|NOFORCE

specifies whether the catalog entry's copy is to be copied from the catalog recovery area even though the copy itself might be inaccurate.

When you specify FORCE, EXPORTRA ignores time stamp and extent mismatches and the object's catalog-entry copy is copied from the volume's catalog recovery area, even though a loss of data integrity might result. Otherwise, when the catalog recovery area's copy is not accurate, it is not copied and the EXPORTRA command terminates with an error message.

Abbreviations: FRC and NFRC

MASTERPW(password)

specifies the master catalog's master password. The master catalog's master password is required when the master catalog is password-protected.

Abbreviation: MRPW

EXPORTRA EXAMPLES

Using EXPORTRA for All Entries on One Volume: Example 1

This example shows the EXPORTRA function for one volume, V0, owned by the VSAM master catalog. All of the entries contained in the catalog recovery area on V0 and the data sets themselves are copied to a SAM data set on another disk volume. All of the data sets reside entirely within the one volume.

```
//EXPORTRA JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//VLVO DD UNIT=3330,VOL=SER=V0,AMP='AMORG',DISP=OLD
//VOLOUT DD UNIT=3330,VOL=SER=333003,DSN=OUT.FILE,
// DISP=(NEW,KEEP),SPACE=(CYL,(5,1)),
// DCB=BLKSIZE=6000
//SYSIN DD *
EXPORTRA -
        OUTFILE(VOLOUT) -
        CRA ( -
        (VLVO ALL))
/*
```

The job control language statements are:

- VLVO DD, which identifies and allocates the volume whose catalog recovery area contents are to be copied. The data sets reside wholly within the one volume.
- VOLOUT DD, which identifies the sequential file that is to receive the copied data sets and catalog recovery records. The DCB BLKSIZE parameter overrides the EXPORTRA default of 2048 to improve performance.

The EXPORTRA command copies everything appearing in the catalog recovery area of volume V0. The EXPORTRA command parameters are:

- OUTFILE, which is required and identifies the sequential (SAM) nonVSAM data set that is to receive the copied information.
- CRA, which is required and identifies the catalog recovery area and volume from which the copy is to take place. VLVO identifies the DD statement for the catalog recovery volume. The ALL subparameter specifies that all entries are to be copied from the catalog recovery area. The INFILE subparameter is not required because the data sets are contained wholly within the one volume.

Using EXPORTRA for All Entries on Multiple Volumes: Example 2

This example shows the EXPORTRA function for one volume owned by the VSAM master catalog. All data sets contained on the volume, V0, and their corresponding CRA entries describing the data sets are copied to a sequential (SAM) data set on a tape volume. The catalog recovery area contains data sets which reside on volumes V1 and V2 in addition to the recovery volume, V0. The entries in the catalog recovery areas of the additional volumes are not to be recovered.

```

//EXPORTRA JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//VLV0 DD UNIT=3330,VOL=SER=V0,AMP='AMORG',
// DISP=OLD
//VLV1 DD UNIT=3330,VOL=SER=V1,AMP='AMORG',
// DISP=OLD
//VLV2 DD UNIT=3330,VOL=SER=V2,AMP='AMORG',
// DISP=OLD
//VOLA DD UNIT=(3330,3),VOL=SER=(V0,V1,V2),
// AMP='AMORG',DISP=OLD
//VOLOUT DD DSN=OUT.FILE,LABEL=(1,SL),
// DISP=(NEW,KEEP),
// UNIT=2400-3,VOL=SER=TAPE01,
// DCB=(BLKSIZE=8192,DEN=3)
//SYSIN DD *
EXPORTRA -
OUTFILE(VOLOUT) -
CRA ( -
(VLV0 ALL INFILE(VOLA)) -
(VLV1 NONE) -
(VLV2 NONE)) -
MASTERPW(MCATMRPW)
/*

```

The job control language statements are:

- VLV0 DD, which identifies the volume whose catalog recovery area contents are to be exported.
- VLV1 and VLV2 DD, each of which identifies a volume containing part of a multivolume data set but whose catalog recovery area contents are not to be exported.
- VOLA DD, which identifies all the volumes containing the data sets to be exported.
- VOLOUT DD, which identifies the sequential tape file that is to receive the exported data sets and catalog recovery records. The DCB BLKSIZE parameter overrides the default of 2048 to improve performance.

The EXPORTRA command copies every entry appearing in the catalog recovery area of volume V0 and the data sets described by the CRA entries. Two additional volumes, V1 and V2, are referenced since they contain part of multivolume data sets, but the entries in their catalog recovery areas are not exported. The EXPORTRA command parameters are:

- OUTFILE, which is required and identifies the sequential (SAM) nonVSAM data set that is to receive the exported information.
- CRA, which is required and identifies the catalog recovery area and volumes from which the copy is to take place. VLV0 identifies the DD statement of the recovery volume. The ALL subparameter specifies that all entries in the catalog recovery area of the volume identified by VLV0 are to be copied. The INFILE subparameter is required, since the data sets to be copied are contained on multiple volumes. VOLA references the DD statement which identifies these multiple volumes, including the recovery volume. The VLV1 NONE and VLV2 NONE subparameters are necessary because the volumes they identify are referenced in a previous INFILE subparameter; however, the contents of their catalog recovery areas are not to be copied.
- MASTERPW, which specifies the master password of the master catalog. This parameter is required when the master catalog is password protected.

Using EXPORTRA for Selected Entries: Example 3

This example shows the EXPORTRA function for selected entries from the catalog recovery area of one volume owned by the VSAM master catalog. Three data sets and their corresponding entries from the catalog recovery area on volume V0 are copied to a sequential (SAM) data set on another disk volume. The remainder of the data sets and the catalog recovery area entries are not copied. Two of the selected data sets are multivolume VSAM data sets. The data sets and their corresponding entries in the catalog recovery areas of the additional volumes are not recovered.

```
//EXPORTRA JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//VLV0 DD UNIT=3330,AMP='AMORG',DISP=OLD,
// VOL=SER=V0
//VLV1 DD UNIT=3330,AMP='AMORG',DISP=OLD,
// VOL=SER=V1
//VLV2 DD UNIT=3330,AMP='AMORG',DISP=OLD,
// VOL=SER=V2
//VOLA DD UNIT=(3330,3),AMP='AMORG',DISP=OLD,
// VOL=SER=(V0,V1,V2)
//VOLB DD UNIT=(3330,2),AMP='AMORG',DISP=OLD,
// VOL=SER=(V0,V1)
//VOLOUT DD UNIT=3330,VOL=SER=333003,DSN=OUT.FILE,
// DISP=(NEW,KEEP),SPACE=(CYL,(5,1))
//SYSIN DD *
EXPORTRA -
  OUTFILE(VOLOUT) -
  CRA ( -
    (VLV0 ENTRIES ( -
      (LARGE.DATASET.A VOLA) -
      (LARGE.DATASET.B VOLB) -
      (LARGE.DATASET.C))) -
    (VLV1 NONE) -
    (VLV2 NONE)) -
  FORCE
/*
```

The job control language statements are:

- VLV0 DD, which identifies the volume whose catalog recovery area contains the entries for the three VSAM data sets to be copied.
- VLV1 and VLV2 DD, each of which identifies a volume containing part of a multivolume data set but whose catalog recovery area contents are not to be copied.
- VOLA DD, which identifies the multivolumes of LARGE.DATASET.A
- VOLB DD, which identifies the multivolumes of LARGE.DATASET.B
- VOLOUT DD, which identifies the sequential file that is to receive the copied data sets and catalog recovery records.

The EXPORTRA command exports the three entries—LARGE.DATASET.A, LARGE.DATASET.B, and LARGE.DATASET.C—from the catalog recovery area on volume V0. Two additional volumes are referenced, since they contain part of multivolume data sets but the entries in their catalog recovery areas are not exported. The EXPORTRA command parameters are:

- OUTFILE, which is required and identifies the nonVSAM data set that is to receive the exported information.
- CRA, which is required and identifies the catalog recovery area and volumes from which the export is to take place. VLV0 identifies the DD statement of the recovery volume. The ENTRIES subparameter identifies the three entries which are

to be copied. LARGE.DATASET.A requires the VOLA subparameter because the VSAM data set is contained on three volumes. Similarly, LARGE.DATASET.B requires the VOLB subparameter. However, LARGE.DATASET.C is contained wholly within the catalog recovery volume, and, therefore, does not require an additional dname3 subparameter. The VLV1 NONE and VLV2 NONE subparameters are necessary because the volumes they identify are referenced in previous dname3 subparameters (VOLA and VOLB); however, the contents of their catalog recovery areas are not to be copied.

- FORCE, which specifies that timestamp mismatches are to be ignored.

IMPORT

The format of the IMPORT command, when it is used to connect a user catalog, is:

IMPORT	CONNECT OBJECTS ((usercatname DEVICETYPE(devtype) VOLUMES(volser))) [CATALOG(mastercatname[/password])]
---------------	---

The format of the IMPORT command, when it is used to move or restore a cluster or alternate index, is:

IMPORT	{INFILE(dname) INDATASET(entryname)} {OUTFILE(dname[/password]) OUTDATASET(entryname[/password])} [ERASE NOERASE] [INTOEMPTY] [OBJECTS((entryname [FILE(dname)] [KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])]) [NEWNAME(newname)] [ORDERED UNORDERED] [VOLUMES(volser[b volser...])] [(entryname...)...])]) [PURGE NOPURGE] [SAVRAC NOSAVRAC] [CATALOG(catname[/password])]
---------------	---

IMPORT may be abbreviated: IMP.

IMPORT PARAMETERS**Required Parameters****CONNECT**

specifies that the entry to be imported is a user catalog. The user catalog is connected to the master catalog in the receiving system. If CONNECT is coded, OBJECTS must also be coded to name the catalog and to identify the volume serial number and device type of the volume that contains the user catalog.

Abbreviation: CON

INFILE(dname)

specifies the name of a DD statement that identifies the portable copy of the cluster or alternate index to be imported. If a non-labeled tape or a direct-access data set created by DOS/VS Access Method Services contains the copy, the following DCB parameters must be specified on the referenced DD statement:

- **BLKSIZE.** If you specified BLKSIZE when the cluster or alternate index was exported, you must specify the same blocksize value for IMPORT. (**Note:** If you did not specify a blocksize for EXPORT, a default value of 2048 was used. Consequently, if you do not specify BLKSIZE for IMPORT, IMPORT sets the blocksize to 2048.)
- **LRECL.** LRECL is based on the maximum record size of the exported VSAM data set. Maximum record size is determined by the value specified via the maximum

subparameter of the RECORDSIZE parameter of the DEFINE CLUSTER or DEFINE ALTERNATEINDEX command when the data set was defined. If the portable data set was created on a system with the OS/VS2 MVS Supervisor Performance #2 Selectable Unit VS2.03.807, LRECL is as follows:

- For relative record data sets, the greater of 276, or maximum VSAM logical record size plus 8
- For all other types of VSAM data sets, the greater of 272, or maximum VSAM logical record size plus 4
- If the portable data set was created on any system without VS2.03.807,
 - For relative record data sets, the greater of 268, or VSAM logical record length plus 8
 - For all other types of VSAM data sets, the greater of 264, or maximum VSAM logical record length plus 4

(Note: If you do not specify LRECL, IMPORT uses a default value of BLKSIZE - 4.)

- RECFM. Must be VBS.

Abbreviation: IFILE

INDATASET(entryname)

names the cluster or alternate index to be imported.

If INDATASET is specified, the entryname is dynamically allocated. The entryname must be cataloged in a catalog that is accessible by the system into which the entry is to be imported. The characteristics of the input data set are described under the INFILE parameter.

Abbreviation: IDS

OUTFILE(dname[/password])

specifies the name of a DD statement that identifies the data set name and volume(s) that is to receive the cluster or alternate index that is to be imported. If the data set was permanently exported and/or you are importing to a volume other than the original volume, the DD statement specifies the name of the cluster or alternate index as DSNAME, the volume serial number(s), the device type, DISP=OLD, and AMP='AMORG'. In addition, you must use concatenated DD statements when both of the following conditions exist and the data set was permanently exported:

- An alternate index or key-sequenced cluster is being imported.
- Its data and index components are on different device types.

The first DD statement specifies the name of the cluster or alternate index as the DSNAME, the volume serial numbers and device type of the data component, DISP=OLD, and AMP='AMORG'. The second DD statement specifies the name of the index component as the DSNAME, the volume serial numbers and device type of its index component, DISP=OLD, and AMP='AMORG'.

If NEWNAME is specified for the cluster or alternate index entry, the data-set name on the DD statement must be the same as the new name. The volume that is to receive the imported cluster or alternate index must be owned by a VSAM catalog.

password

is the update (or higher-level) password of the output cluster or alternate index. You must supply the password when the output data set is empty (see the VOLUMES parameter for more information about importing into an empty data set). Otherwise, IMPORT obtains the required VSAM data set passwords from the portable data set.

When you use a DD statement to describe the output data set, you must identify the data set's catalog with a JOBCAT or STEPCAT DD statement unless:

- The data set's entry is in the master catalog, or
- The first qualifier of the data set's qualified name identifies a catalog (that is, the first qualifier is the catalog's name or alias).

Note: When importing a cluster that was permanently exported, the OUTFILE parameter should be used.

Abbreviation: OFILE

OUTDATASET(entryname[/password])

specifies the name of the data set that is to receive the data being imported. When you specify OUTDATASET, the VSAM data set you identify is dynamically allocated. If NEWNAME is specified for the cluster or alternate index entry, entryname must be the same as the new name.

password

is the update- (or higher-level) password of the output object. See password under INFILE for password requirements of the output data set.

Abbreviation: ODS

Optional Parameters**CATALOG(catname[/password])**

specifies the name of the catalog in which the imported object is to be cataloged. This parameter is required when the catalog is password protected.

catname

is the name of the catalog in which the entry to be imported is to be defined. If you are importing a user catalog, the specified catalog must be the master catalog.

password

specifies the catalog's update or higher level password. When you import an alternate index whose base cluster is password-protected, you should supply the catalog's master password. Otherwise, you (or the operator) will be prompted to supply the base cluster's master password.

Abbreviation: CAT

ERASE|NOERASE

specifies whether the data component of the cluster or alternate index that was exported with the TEMPORARY option is to be erased (that is, overwritten with binary zeros). This parameter can be used only when you are importing the object into the system from which it was previously exported with the TEMPORARY option. This parameter overrides whatever was specified when the object was defined or last altered.

Abbreviations: ERAS and NERAS

INTOEMPTY

specifies that you are importing into an empty data set. If this parameter is not specified, an attempt to import into an empty data set will fail. Importing into an empty data set causes the data set to be loaded from the portable data set. The password and RACF profiles associated with the empty data set will be retained. When importing into an empty data set, the SAVRAC | NOSAVRAC parameter applies only to the paths imported and successfully defined over the empty data set. If the define of an exported path fails because a catalog entry with the same name already exists, the path on the portable data set is ignored.

Abbreviation: IEMPTY

OBJECTS

specifies the new or changed attributes for the cluster, alternate index, any associated paths, or user catalog to be imported.

By specifying the OBJECTS parameter:

- You provide the information needed to connect a user catalog to the master catalog, or
- You may override certain attributes of the cluster or alternate index (and associated paths) to be imported in order to provide new or changed attributes. Access Method Services matches each entryname you specify through the OBJECTS parameter against the name of each object on the portable data set. When a match is found, the information specified by OBJECTS overrides the information on the portable data set. The description for each of the OBJECTS subparameter includes information on how they should be specified; that is, what name you should specify as the entryname to accomplish the results you want.

You can specify OBJECTS

- For a user catalog:

```
OBJECTS((usercatname
        DEVICETYPE(devtype)
        VOLUMES(volser)))
```

- For a cluster, alternate index, or path:

```
OBJECTS((entryname
        [FILE(dname)]
        [KEYRANGES((lowkey b highkey)
                  [(lowkey b highkey)...])]
        [NEWNAME(newname)]
        [ORDERED|UNORDERED]
        [VOLUMES(volser[b volser...])]
        [(entryname...)...])
```

usercatname
specifies the name of the user catalog being connected.

entryname
specifies the name of the data component, index component, cluster, alternate index, or path for which attributes are being specified. For a cluster, alternate index, data component, index component, or path, this name must appear on the portable data set;

otherwise, the parameter list will be ignored.

Abbreviation: OBJ

DEVICETYPE(devtype)

specifies the device type of the volume that contains a user catalog that is to be connected. You can specify a device type for any direct-access device that is supported by your OS/VS2 MVS system.

Abbreviation: DEVT

FILE(dname)

specifies the name of a DD statement that identifies the volumes allocated to the data and index components of a key-sequenced cluster or of an alternate index. This parameter is required when the components are defined as unique and when the data and index components reside on different device types. When components reside on different device types, FILE must be coded twice within the OBJECTS parameter: once in the parameter set for the index component and once in a second parameter set for the data component. If you do not specify FILE, the required volumes are dynamically allocated.

KEYRANGES([lowkey b highkey][[lowkey b highkey)...])

specifies portions of a key-sequenced cluster or alternate index to be placed on separate volumes.

The data is divided, by key, among the volumes as contained in the portable data set or as specified in the VOLUMES parameter. If a volume serial number is duplicated on the portable data set or in VOLUMES, multiple key ranges of an alternate index or cluster attribute are placed on that volume. If the number of volumes is greater than the number of key ranges, the excess volumes are used for overflow records from any key range without consideration of range boundaries. If there are fewer volumes than key ranges, the excess key ranges are placed on the last volume specified.

By using the KEYRANGE parameter you may specify key range pairs for a key-sequenced data set that previously was not divided by key or that was divided in different key range groups.

The maximum number of key-range pairs is 20. Key ranges must be in ascending order, and may not overlap. Gaps may exist within a specified set of ranges, but records cannot be inserted within a gap.

When you specify KEYRANGES, entryname can be either the cluster or alternate index name or the name of the data component. If you specify KEYRANGES with the cluster or alternate index name, the values specified are propagated to and defined for the data component. If you specify KEYRANGES with the data component name the values are defined for the data component and any specification of KEYRANGES with the cluster or alternate index name is overridden.

Keys can contain 1 to 64 characters; if coded in hexadecimal, they may contain 1 to 128 hexadecimal characters. All EBCDIC characters are allowed. Keys consisting of characters must be enclosed in single quotation marks if they contain commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in single quotation marks. If the key is specified in hexadecimal, it must be preceded by X and be enclosed in single quotation marks.

lowkey

specifies the low key of the key range. If lowkey is shorter than the actual keys, it will be padded with binary zeros.

highkey

specifies the high key of the key range. If highkey is shorter than the actual keys, it will be padded with binary ones.

Abbreviation: KRNG

NEWNAME(newname)

specifies the new name of an imported cluster or alternate index or its components, or an associated path. When you specify NEWNAME, only the name specified as entryname is changed. You cannot specify NEWNAME if entryname identifies a user catalog.

The new name can contain 1 to 44 alphanumeric characters (A through Z, and 0 through 9), national characters (à, #, and \$), and special characters (the hyphen and 12-0 overpunch). Names that contain more than eight characters must be segmented by periods; one to eight characters may be specified between periods. The first character of the name or name segment must be either an alphabetic character or a national character. If you specify the new name in the format newname(modifier), only that portion of the name preceding the left parenthesis will be used! The modifier enclosed in parentheses will be ignored.

If you are specifying a new name for a cluster or alternate index that was exported with the TEMPORARY option and it is being imported back into the original system, you must rename each of its components, also.

Abbreviation: NEWNM

ORDERED|UNORDERED

specifies whether volumes are to be used in the order in which they were listed when the data set was originally defined (as contained on the portable data set) or in the VOLUMES parameter. If the data set is divided into key ranges, all of the records within the range specified by the first low-key/high-key pair are placed on the first volume; all of the records within the second range are placed on the second volume; etc. If it is impossible to allocate volumes in the given order and ORDERED is specified, the command is terminated.

When you specify ORDERED|UNORDERED, you can specify the cluster or alternate index name, the data component name or the index component name as entryname with the following results:

- If ORDERED|UNORDERED is specified with the cluster or alternate index name, the specified attribute is propagated to and defined for the data and, for a key-sequenced cluster or alternate index, the index component.
- If ORDERED|UNORDERED is specified with the data component name, the specified attribute is defined for the data component. Any specification of ORDERED|UNORDERED with the cluster or alternate index name is overridden.
- For a key-sequenced cluster and for an alternate index, if ORDERED|UNORDERED is specified with the index component name, the specified attribute is

defined for the index component. Any specification of ORDERED|UNORDERED with the cluster or alternate index name is overridden.

Abbreviations: ORD and UNORD

VOLUMES(volser[b volser...])

specifies the volumes on which the cluster or alternate index is to reside or the volume on which the user catalog resides.

The volume serial number may contain one to six alphameric, national (@, #, and \$), hyphens, and special (commas, semicolons, blanks, parentheses, slashes, asterisks, periods, quotation marks, ampersands, plus signs, and equal signs) characters. The volume serial numbers must be enclosed in single quotation marks if they contain special characters. Single quotation marks within a volume serial number must be coded as two single quotation marks.

If VOLUMES is not coded, the original volume is the receiving volume. This parameter is required when a user catalog is to be imported; when importing a user catalog, specify only one volume. If the portable data set was created with an EXPORT command using the PERMANENT option on a Release 1 version of Access Method Services (DOS/VS Release 29, OS/VS1 Release 2, OS/VS2 Release 1.6, or OS/VS2 Release 2), this parameter is required. Portable data sets created by Release 1 using the PERMANENT option do not contain volume information.

When you specify VOLUMES, you can specify the cluster or alternate index name, the data component name or the index component name as entryname with the following results:

- If VOLUMES is specified with the cluster or alternate index name, the specified volume list is propagated to and defined for the data and, for a key-sequenced cluster or alternate index, the index component.
- If VOLUMES is specified with the data component name, the specified volume list is defined for the data component. Any specification of VOLUMES with the cluster or alternate index name is overridden.
- For a key-sequenced cluster or alternate index, if VOLUMES is specified with the index component name, the specified volume list is defined for the index component. Any specification of VOLUMES with the cluster or alternate index name is overridden.

For clusters or alternate indexes, multiple volumes, if specified, must be of the same device type. By repeating the OBJECTS parameter set for each component and including VOLUMES in each parameter set, you can have the data and index components on different volumes. Although the index and data components may reside on different device types, each volume of a multivolume component must be of the same type.

If the receiving volume is of a type different from that that originally contained the cluster or alternate index, the job may be terminated because of allocation problems. Each space allocation quantity is recorded in a catalog entry as an amount of cylinders or tracks even when RECORDS was specified in the DEFINE command. When a cluster or alternate index is imported, the number of cylinders or tracks in the catalog entry is not modified, even though the object may be imported to reside on a device type other than

that it was exported from. An attempt to import an object that previously resided on a 3330 may fail if it is imported to a 2314. Conversely, if an object is exported from a 2314 and imported to a 3330, more space is allocated than the object needs.

You can avoid space allocation problems when you define an empty cluster or alternate index on the new system or catalog and identify it as the target for the object being imported as described below:

- Using the DEFINE command, define a new entry for the cluster or alternate index in the catalog to which it is to be moved. If space was allocated in RECORDS, you may specify the same quantity; if it was allocated in TRACKS or CYLINDERS, you must adjust the quantity for the new device type. If an entry already exists in the catalog for the object, you must delete that entry or use a different name in the DEFINE command.
- Using the IMPORT command, load the portable data set into the newly defined cluster or alternate index. When IMPORT encounters an empty target data set, the exported catalog information is bypassed and just the data records are processed.

Abbreviation: VOL

PURGE|NOPURGE

specifies whether the original cluster or alternate index is to be deleted and replaced regardless of the retention time specified in the TO or FOR parameter. This parameter can be used only when you are importing the object into the original system from which it was exported with the TEMPORARY option.

Abbreviations: PRG and NPRG

SAVRAC|NOSAVRAC

specifies, for a RACF-protected object, whether existing profiles are to be used or whether new profiles are to be created.

SAVRAC should be specified when RACF data set profiles already exist for objects being imported from the portable data set. Typically, you would specify this option when replacing a data set with a portable copy made with an EXPORT TEMPORARY operation. SAVRAC will cause the existing profiles to be "saved" and used, rather than letting the system delete old profiles and create new, default profiles.

Caution: You should ensure that valid profiles do exist for all components (for example, cluster, data, index) being imported when SAVRAC is specified. If this is not done, an invalid and possibly improper profile may be "saved" and used inappropriately. Remember, that paths are imported along with their corresponding cluster or alternate index, and the same caution applies to these entries. In particular, keep in mind that additional paths may be brought in during the import.

NOSAVRAC should be specified when you wish new profiles to be created. This is usually the situation when importing a permanently exported cluster or alternate index. A profile will be defined for the imported components if either the Automatic Data Set Protection option has been specified for you or if the exported component had a RACF indication in the catalog when it was exported. If you import into a catalog in which there is a component with a duplicate name which is marked as having been temporarily exported, it and any associated profiles will be deleted before importing the portable data set.

IMPORT EXAMPLES

Importing a User Catalog: Example 1

In this example, a user catalog, D27UCAT1, is imported and connected to the new system's master catalog, AMASTCAT. This example reconnects the user catalog, D27UCAT1, that was disconnected with EXPORT Example 3.

```
//IMPORT1 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
IMPORT -
  OBJECTS( -
    (D27UCAT1 -
      VOLUME(VSER02) -
      DEVICETYPE(2314) ) -
    ) -
  CONNECT -
  CATALOG(AMASTCAT/MRCATPW2)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The IMPORT command builds a user catalog connector entry that identifies the user catalog D27UCAT1 in the master catalog AMASTCAT. Its parameters are:

- OBJECTS, which is required when a user catalog is being imported. The subparameters of OBJECTS identify the user catalog, D27UCAT1, the user catalog's volume, VSER02, and the device type of the user catalog's volume, 2314.
- CONNECT, which specifies that the user catalog connector entry is to be built and put in the master catalog to connect the user catalog to the master catalog. CONNECT is required when a user catalog is being imported.
- CATALOG, which identifies the master catalog, AMASTCAT, and specifies its update- (or higher-level) password, MRCATPW2.

Importing a Key-Sequenced Cluster: Example 2

In this example, a key-sequenced cluster, D40.EXAMPLE.KSDS1, that was previously exported, is imported. (See the EXPORT example, "Exporting a Key-Sequenced Cluster.") OUTFILE and its associated DD statement are provided to allocate the data set.

The original copy of D40.EXAMPLE.KSDS1 is replaced with the imported copy, TAPE2. Access Method Services finds and deletes the duplicate name, D40.EXAMPLE.KSDS1, in the catalog D27UCAT1. (A duplicate name exists because TEMPORARY was specified when the cluster was exported.) Access Method Services then redefines D40.EXAMPLE.KSDS1 using the catalog information from the portable file, TAPE2.

```

//IMPORT2 JOB      ...
//STEP1  EXEC      PGM=IDCAMS
//SOURCE DD        DSNAME=TAPE2,UNIT=(2400-3,,DEFER),
//      VOL=SER=003030,DISP=OLD,
//      DCB=(BLKSIZE=6000,LRECL=479,DEN=3),LABEL=(1,SL)
//SYSPRINT DD      SYSOUT=A
//SYSIN   DD        *
      IMPORT -
      INFILE(SOURCE) -
      OUTDATASET(D40.EXAMPLE.KSDS1) -
      CATALOG(D27UCAT1/MRPWD27)
/*

```

The job control language statements are:

- SOURCE DD, which describes the portable data set, TAPE2. TAPE2 resides on a magnetic tape file, which will not be mounted by the operator until Access Method Services opens TAPE2 for processing. The blocksize parameter is included (even though it need not be, because the tape has a standard label and the information is contained in the data-set header label) to illustrate the fact that the information specified when the data set is imported is required to be the same as was specified when the data set was exported. The LRECL parameter is not required, because the maximum record size is 475 bytes and the default (blocksize minus 4) is adequate. However, by specifying a recordsize, the default is overridden and virtual storage is more efficiently used. When recordsize is specified, it is the largest-record size + 4.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The IMPORT command copies the portable data set, TAPE2, into the system and assigns it the name D40.EXAMPLE.KSDS1. When TAPE2 is copied, Access Method Services reorganizes the data records so that deleted records are removed and control intervals and control areas contain the specified freespace percentages. The original copy of the cluster is deleted and replaced with the data records from the TAPE2 portable file. The IMPORT command's parameters are:

- INFILE, which points to the SOURCE DD statement. The SOURCE DD statement describes the portable file, TAPE2, to be imported.
- OUTDATASET, which gives the name of the data set being imported. Since the high-level qualifier of the data set is the name of the alias of the user catalog D27UCAT2, Access Method Services can dynamically allocate the cluster without specifying a JOBCAT or STEPCAT DD statement.
- CATALOG, which identifies the catalog, D27UCAT1, in which the imported cluster is to be defined. The catalog's update- (or higher-level) password is required.

Importing an Entry-Sequenced VSAM Cluster: Example 3

In this example, an entry-sequenced cluster, D50.EXAMPLE.ESDS1, is imported from a portable file, TAPE1. This example is associated with EXPORT Example 2, "Exporting an Entry-Sequenced Cluster." The cluster is defined in a different catalog from which it was exported, assigned a new name, and imported to a different volume.

```

//IMPORT3 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SOURCE DD DSN=TAPE1,UNIT=(2400-3,,DEFER),DISP=OLD,
// VOL=SER=001147,LABEL=(1,SL)
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
IMPORT -
  INFILE(SOURCE) -
  OUTDATASET(D40.EXAMPLE.ESDS3) -
  OBJECTS( -
    (D50.EXAMPLE.ESDS1 -
      NEWNAME(D40.EXAMPLE.ESDS3) -
      VOLUMES(VSER02) ) -
    ) -
  CATALOG(D27UCAT1/UPPWD27)
/*

```

The job control language statements are:

- SOURCE DD, which describes the portable file, TAPE1. TAPE1 resides on a magnetic tape file, which will not be mounted by the operator until Access Method Services opens TAPE1 for processing.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The IMPORT command moves the contents of the portable file, TAPE1, into the system. When TAPE1 is moved, Access Method Services reorganizes the data records. The IMPORT command's parameters are:

- INFILE, which points to the SOURCE DD statement. The SOURCE DD statement describes the data set to be imported, TAPE1.
- OUTDATASET gives the name of the data set that is being imported—the name is the renamed cluster. Since the high-level qualifier of the data set name is the name of the alias of the catalog, D27UCAT1, the data set can be dynamically allocated without specifying a JOBCAT or STEPCAT DD statement.
- OBJECTS, which changes some of the attributes for the object being imported:
 - D50.EXAMPLE.ESDS1, which identifies the entry-sequenced cluster as it is currently named on TAPE1.
 - NEWNAME, which specifies that the cluster's entryname is to be changed to D40.EXAMPLE.ESDS3.
 - VOLUMES, which identifies the new volume on which the cluster is to reside.
- CATALOG, which identifies the catalog, D27UCAT1, that is to contain the cluster's catalog entry. The catalog's update- (or higher-level) password is required.

IMPORTRA

The format of the IMPORTRA command is:

IMPORTRA	{INFILE(dname) INDATASET(entryname)} [OUTFILE(dname)] OBJECTS((entryname [DEVICETYPES(devtype[b devtype...])] [FILE(dname)] [VOLUMES(volser[b volser...])])] [(entryname...)]...)] [SAVRAC NOSAVRAC] [CATALOG(catname[/password])]
-----------------	--

IMPORTRA can be abbreviated: MPRA.

IMPORTRA PARAMETERS

Required Parameters

INFILE(dname)

names the DD statement that describes the portable data set (that is, the data set that resulted when you issued the EXPORTRA command). If the portable data set resides on a nonlabeled tape or in a direct-access data set created by DOS/VS Access Method Services, the DCB BLKSIZE parameter must be specified on the referenced DD statement. If you specified BLKSIZE for the portable data set which was created when you executed your EXPORTRA job, you must specify the same blocksize for your IMPORTRA job. (**Note:** If you did not specify a blocksize for EXPORTRA, a default value of 2048 was used. Consequently, if you do not specify BLKSIZE for IMPORTRA, IMPORTRA sets the blocksize to 2048.) All other characteristics of the portable data set are set by IMPORTRA.

Abbreviation: IFILE

INDATASET(entryname)

specifies the data set name of the portable data set. When you specify INDATASET instead of INFILE, the portable data set (that is, the data set that resulted when you issued the EXPORTRA command) is dynamically allocated.

Abbreviation: IDS

Optional Parameters

CATALOG(catname[/password])

identifies the target catalog (that is, the catalog that is to contain the imported catalog entries).

If you do not include the CATALOG parameter, the JOBCAT or STEPCAT catalogs are used. If you have not specified a JOBCAT or STEPCAT catalog, the VSAM master catalog is used. You must specify the CATALOG parameter when the target catalog is password protected. Otherwise, the system operator is prompted to supply the correct password.

catname

names the catalog. When you specify a user catalog, you must describe and allocate the catalog with a JOBCAT or STEPCAT DD statement.

password

is the target catalog's update- or higher-level password, if the catalog is password protected. If any of the data sets being imported is an alternate index whose base cluster is password protected, you should supply the catalog's master password.

Abbreviation: CAT

```
OBJECTS((entryname
  [DEVICETYPES(devtype[b devtype...])]
  [FILE(dname)]
  [VOLUMES(volser[b volser...])])
  [(entryname...)]...)]
```

The OBJECTS parameter group specifies new or changed attributes for objects on the portable data set.

Attributes may be specified for a cluster, alternate index, user catalog, or nonVSAM data set. By specifying the OBJECTS parameter, you may override certain attributes contained on the portable data set. Access Method Services matches each entryname you specify against the name of each object on the portable data set. When a match is found, the information specified by OBJECTS overrides the information on the portable data set.

entryname

specifies the object's entryname. You can specify the entryname and associated attributes for up to 255 objects.

You can specify the entryname of these types of objects only: VSAM cluster or alternate index, nonVSAM data set, user catalog, and a VSAM object's data or index component.

Abbreviation: OBJ

```
DEVICETYPES(devtype[b devtype...])
```

specifies the device type on which a user catalog being imported resides or the device type(s) on which a nonVSAM data set being imported resides. If entryname names a user catalog, you may specify only one device type. If entryname names a nonVSAM data set which resides on different device types, the device types listed in the DEVICETYPES parameter must be specified in the same order as the volume serial numbers listed in the VOLUMES parameter.

This parameter can be specified only when the entryname names a nonVSAM data set or user catalog. For a nonVSAM data set or user catalog, if DEVICETYPES is not coded, the data set or user catalog must still reside on the same device type(s) as contained in the entry in the portable data set.

Abbreviation: DEVT

```
FILE(dname)
```

specifies the name of a DD statement that identifies the volumes allocated to the data and index components of a key-sequenced cluster or of an alternate index. This parameter is required when the components are defined as unique and when the data and index components reside on different device types. When components reside on different device types, FILE must be coded twice within the OBJECTS parameter; once in the parameter set for the index component and once in a second parameter set for the data component. If you do not specify FILE, the required volumes are dynamically allocated.

VOLUMES(volser[b volser...])

specifies the volume(s) on which a cluster or alternate index is to reside, the volume on which a user catalog resides, or the volume(s) on which a nonVSAM data set resides. If entryname names a user catalog, you may specify only one volume serial number as a subparameter of VOLUMES.

For a cluster or alternate index, if VOLUMES is not coded, the original volume is the receiving volume. For a user catalog or nonVSAM data set, if VOLUMES is not coded, the data set or user catalog must still reside on the same volume(s) as contained in the portable data set.

volser

is a 1 to 6 alphameric or special character volume serial number.

If the password contains a lowercase alphabetic character, it will be folded to an uppercase alphabetic character.

When the volser contains a special character (that is, a comma [,], semicolon [;], blank [b], period [.] , single quote ['], ampersand [&], plus sign [+], equal sign [=], hyphen [-], parenthesis [(or)], slash [/], or asterisk [*]), enclose the volser in single-quotation marks (for example, VOLUMES('*DORIS')).

When the volser contains a special character and also contains a single-quotation mark, code the imbedded quotation mark as two single-quotation marks (for example, VOLUMES('*CA''RO')).

You can code volser in hexadecimal form, where two hexadecimal characters represent one alphameric or special character. For example, VOLUMES(X'E2E4C5') is the same as VOLUMES(SUE).

When you specify VOLUMES for a cluster or alternate index, you can specify the cluster or alternate index name, the data component name or the index component name as entryname with the following results:

- If VOLUMES is specified with the cluster or alternate index, name, the specified volume list is propagated to and defined for the data and, for a key-sequenced cluster or alternate index, the index component.
- If VOLUMES is specified with the data component name, the specified volume list is defined for the data component. Any specification of VOLUMES with the cluster or alternate index name is overridden.
- For a key-sequenced cluster or alternate index, if VOLUMES is specified with the index component name, the specified volume list is defined for the index component. Any specification of VOLUMES with the cluster or alternate index name is overridden.

For cluster or alternate indexes, if multiple volumes are specified, they must be of the same device type. By repeating the OBJECTS parameter set for each component and including VOLUMES in each parameter set, you can have the data and index components on different volumes. (See also the description of the FILE parameter.) Although the index and data components may reside on

different device types, each volume of a multivolume component must be of the same type.

Abbreviation: VOL

OUTFILE(dname)

names the DD statement that contains a data set name and the volume serial number of each volume that is to contain the imported VSAM clusters and alternate indexes. You must use concatenated DD statements if the data sets are on different device types. If the OUTFILE parameter is not supplied, the required data sets are dynamically allocated as needed.

The dsname specified with the DD statement cannot be one of the names cataloged in the target catalog (that is, the catalog that is to contain the imported catalog entries) and it cannot be one of the entrynames within the portable data set. The section "Restoring the Catalog Entry that was Obtained Using the EXPORTRA Command" describes how this dsname is used by IMPORTRA processing. The DD statement(s) also specify the volume serial number(s), device type, DISP=OLD, and AMP='AMORG'.

Abbreviation: OFILE

SAVRAC|NOSAVRAC

specifies, for a RACF-protected object, whether existing profiles are to be used or whether new profiles are to be created.

SAVRAC should be specified when RACF data set profiles already exist for objects being imported from the portable data set. SAVRAC will cause the existing profiles to be "saved" and used, rather than letting the system delete old profiles and create new default profiles.

Caution: You should ensure that valid profiles do exist for all components (for example, cluster, data, index) being imported when SAVRAC is specified. If this is not done, an invalid and possibly improper profile may be "saved" and used inappropriately. Remember, that paths are imported along with their corresponding cluster or alternate index, and the same caution applies to these entries. In particular, keep in mind that additional paths may be brought in during the import.

NOSAVRAC should be specified when you wish new profiles to be created. A profile will be defined for the imported components if either the RACF Automatic Data Set Protection option has been specified for you or if the exported component had a RACF indication in the catalog at the time it was exported. The mode, SAVRAC or NOSAVRAC, set by these keywords applies to all VSAM components being imported back into the system.

There may be situations where the SAVRAC or NOSAVRAC mode would not be suitable for all entities restored from the catalog recovery area. In such cases, the following procedure is suggested:

- Issue two separate EXPORTRA commands, resulting in two portable data sets. One EXPORTRA will create a portable data set containing entities to which SAVRAC should apply, the other to which NOSAVRAC should apply.
- Issue two IMPORTRA commands, one specifying SAVRAC and the other NOSAVRAC.

IMPORTRA EXAMPLE

This example shows how EXPORTRA and IMPORTRA can be used to recover a VSAM catalog. The first part illustrates the use of EXPORTRA. The second part shows how IMPORTRA is used.

Recovering a VSAM Catalog: Example 1—Part 1 (EXPORTRA)

This example performs the EXPORTRA function against the VSAM master catalog, AMASTCAT, and against all of the volumes owned by it. All of the data sets listed in the catalog recovery areas on the catalog volume and on the other volumes owned by AMASTCAT are exported to a SAM data set on another disk volume. Using the FORCE option causes EXPORTRA to ignore time stamp mismatches between the volumes and the catalog.

```
//RECOVER JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CRAVOL1 DD DISP=OLD,VOL=SER=VSER01,UNIT=2314,
// AMP='AMORG'
//CRAVOL2 DD DISP=OLD,VOL=SER=VSER04,UNIT=3330,
// AMP='AMORG'
//PORT DD DSNAME=PORT,DISP=(NEW,KEEP),
// SPACE=(CYL,(3,3)),VOL=SER=231401,
// UNIT=2314,DCB=BLKSIZE=8000
//SYSIN DD *
```

```
EXPORTRA -
  CRA( -
    (CRAVOL1 ALL) -
    (CRAVOL2 ALL) -
  ) -
  FORCE -
  OUTFILE(PORT) -
  MASTERPW(MCATMRPW)
```

/*

The job control language statements are:

- CRAVOL1 DD, which identifies and allocates the first volume whose catalog recovery area contents is to be exported.
- CRAVOL2 DD, which identifies and allocates the second volume whose catalog recovery area contents is to be exported.
- PORT DD, which identifies the sequential file that is to receive the exported data sets and catalog recovery records. The DCB BLKSIZE parameter overrides the EXPORTRA default of 2048 to improve performance.

The EXPORTRA command exports everything appearing in the catalog recovery areas of volumes VSER01 and VSER04. The EXPORTRA command's parameters are:

- CRA, which is required and identifies the catalog recovery areas and volumes from which the export is to take place. The dnames of the DD statements for these objects must be identical to the names specified for this parameter. The ALL subparameter specifies that everything is to be exported from each catalog recovery area.
- FORCE, which specifies that timestamp mismatches are to be ignored.
- OUTFILE, which is required and identifies the sequential (SAM) nonVSAM data set that is to receive the exported information. The dname of the DD statement for this object must be identical to the name specified with this parameter.

- MASTERPW, which specifies the master password of the master catalog. This parameter is required in order to export information from catalog recovery areas controlled by a recoverable catalog.

Recovering a VSAM Catalog: Example 1—Part 2 (IMPORTRA)

This example imports all of the data sets that were exported using EXPORTRA in the previous example. The receiving catalog is the VSAM master catalog, and the CATALOG parameter is used to supply its master password.

```
//RESTORE JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//PORT DD DSN=PORT,DISP=OLD,UNIT=2314,
// VOL=SER=231401,DCB=BLKSIZE=8000
//VSAMIN DD DSN=DUMMY.NAME,DISP=OLD,UNIT=2314,
// VOL=SER=VSER01,AMP='AMORG'
// DD VOL=SER=VSER04,UNIT=3330,DISP=OLD
//SYSIN DD *
```

```
IMPORTRA -
  INFILE(PORT) -
  OUTFILE(VSAMIN) -
  CATALOG(AMASTCAT/MCATMRPW)
```

```
/*
```

The job control language statements are:

- PORT DD, which identifies and allocates the portable data set created previously by EXPORTRA. The DCB BLKSIZE parameter is included (even though it need not be because the information is contained in the DSCB for the data set) to illustrate the fact that the blocksize specified in the IMPORTRA job must be the same as that specified in the EXPORTRA job which created the portable data set.
- VSAMIN DD, which identifies and allocates the volumes of the objects to be imported. This permits them to be reloaded from the portable data set. The dsname appearing on the DD statement is required. It is a dummy name that must not appear either in the catalog or among the names on the portable data set. Because the volumes are of different device types, concatenated DD statements must be used.

The IMPORTRA command imports the data sets written on the portable data set by EXPORTRA previously. The IMPORTRA command's parameters are:

- INFILE, which is required and which identifies the portable data set.
- OUTFILE, which identifies each VSAM volume involved in the import. The DD statement identified by OUTFILE also identifies the dummy data set, DUMMY.NAME.
- CATALOG, which specifies the name of the master catalog. The master password of the target catalog is required in order to import catalog information.

LISTCAT

The format of the LISTCAT command is:

LISTCAT	[CATALOG(catname[/password])] [ENTRIES(entryname[/password] [b entryname[/password]...])] LEVEL(level)] [ALIAS] [b ALTERNATEINDEX] [b CLUSTER] [b DATA] [b GENERATIONDATAGROUP] [b INDEX] [b NONVSAM] [b PAGESPACE] [b PATH] [b SPACE] [b USERCATALOG] [ALL] ALLOCATION HISTORY NAME VOLUME] [CREATION(days)] [EXPIRATION(days)] [NOTUSABLE] [OUTFILE(dname)]
---------	---

LISTCAT may be abbreviated: LISTC.

LISTCAT PARAMETERS

Required Parameters

The LISTCAT command has no required parameters.

The LISTCAT command is issued as a job step (that is, not through TSO) and no parameters are specified; an entire catalog is listed. See the section "Order of Catalog Use: LISTCAT" for a description of how the catalog to be listed is selected.

Note to TSO users: When LISTCAT is invoked from a TSO terminal and no operands are specified, the prefix of the TSO user becomes the highest level of entryname qualification and only those entries with a matching highest level of qualification are listed. It is as if you specified:

LISTCATLEVEL(TSO-user-prefix)

Optional Parameters

[ALIAS][b ALTERNATEINDEX][b CLUSTER][b DATA]
[b GENERATIONDATAGROUP][b INDEX][b NONVSAM]
[b PAGESPACE][b PATH][b SPACE][b USERCATALOG]

specifies that certain types of entries are to be listed. Only those entries whose type is specified is listed (that is, when you specify CLUSTER but not DATA or INDEX, the cluster's entry is listed and its associated data and index entries are not listed). When you identify an entry with its name (that is, when you specify ENTRIES) and also specify an entry type, the names entry is not listed unless it is of the specified type. You can specify as many entry types as desired. When you want to completely list a catalog, do not specify any entry type.

ALIAS

specifies that alias entries are to be listed.

ALTERNATEINDEX

specifies that entries for alternate indexes are listed. If ALTERNATEINDEX is specified and DATA and INDEX are not also specified, entries for the alternate-index's data and index components aren't listed.

Abbreviation: AIX

CLUSTER

specifies that cluster entries are to be listed. If CLUSTER is specified and DATA and INDEX are not also specified, entries for the cluster's data and index components are not listed.

Abbreviation: CL

DATA

specifies that entries for data components of clusters and alternate indexes are to be listed. If a VSAM object's name is specified and DATA is coded, only the object's data-component entry is listed. When DATA is the only entrytype parameter coded, the catalog's data component is not listed.

GENERATIONDATAGROUP

specifies that entries for generation data groups are to be listed.

Abbreviation: GDG

INDEX

specifies that entries for index components of key-sequenced clusters and alternate indexes are to be listed. If a VSAM object's name is specified and INDEX is coded, only the object's index-component entry is listed. When INDEX is the only entrytype parameter coded, the catalog's index component is not listed.

Abbreviation: IX

NONVSAM

specifies that entries for nonVSAM data sets are to be listed. If a generation data group's name and nonVSAM are specified, the generation data sets associated with the generation data group are listed.

Abbreviation: NVSAM

PAGESPACE

specifies that entries for page spaces are to be listed.

Abbreviation: PGSPC

PATH

specifies that entries for paths are listed.

SPACE

specifies that entries for volumes containing data spaces defined in this catalog are to be listed. Candidate volumes are included. If entries are identified by entryname, SPACE can be coded only when no other entrytype parameter is coded.

Abbreviation: SPC

USERCATALOG

specifies that catalog connectors are to be listed. The user-catalog connector entries are in the master catalog. (User-catalog connector entries can also be

in a user catalog, but the OS/VS2 system does not recognize them when searching for a user catalog.)

Abbreviation: UCAT

ALL|ALLOCATION|HISTORY|NAME|VOLUME

specifies the fields to be included for each entry listed. "Appendix B: Interpreting LISTCAT Output Listings" shows the listed information that results when you specify nothing (which defaults to NAME), ALL, VOLUME, ALLOCATION, and HISTORY.

ALL

specifies that all fields are to be listed.

ALLOCATION

specifies that the information provided by specifying VOLUME and detailed information about the allocation are to be listed. The information about allocation is listed only for data and index component entries.

Abbreviation: ALLOC

HISTORY

specifies that only the following information is to be listed for each entry: name, entry type, ownerid, creation date, expiration date, and for a recoverable catalog's entries, the catalog recovery area's volume, device type, and control interval number. It can be specified for CLUSTER, DATA, INDEX, ALTERNATEINDEX, PATH, GENERATIONDATAGROUP, PAGESPACE, and NONVSAM.

Abbreviation: HIST

NAME

specifies that the name and entry type of the entries are to be listed.

Note to TSO users: Only the name of each entry associated with the TSO user's prefix is listed when no other parameters are coded.

Some entrytypes are listed along with their associated entries. The entrytype and name of the associated entry follows the listed entry's name. For details, see "ASN: Associations Group" in "Appendix B: Interpreting LISTCAT Output Listings."

VOLUME

specifies that the information provided by specifying HISTORY, and the volume serial numbers and device types allocated to the entries are to be listed. Volume information is only listed for data and index component entries, data-space (volume) entries, nonVSAM data-set entries, and user-catalog connector entries.

Note to TSO users: Only the name and volume serial numbers associated with the TSO user's prefix are listed when no other parameters are coded.

Abbreviation: VOL

CATALOG[catname[/password]]

specifies the name of the catalog that contains the entries that are to be listed. If CATALOG is coded, only entries from that catalog are listed. See "Order of Catalog Use: LISTCAT" for information about the order in which catalogs are searched.

catname

is the name of the catalog.

password

specifies the read-level or higher-level password of the catalog that contains entries to be listed. If the entries to be listed are password protected, a password must be supplied either through this parameter or through the ENTRIES parameter. If passwords are to be listed, you must specify the master password.

If the catalog's volume is physically mounted, it is dynamically allocated. The volume must be mounted as permanently resident or reserved.

Abbreviation: CAT

CREATION(days)

specifies that entries of the indicated type (CLUSTER, DATA, etc.) are to be listed only if they were created the specified number of days ago or earlier.

days

specifies the number of days ago. This value can be expressed in decimal, hexadecimal, or binary. If it is expressed in hexadecimal or binary, it must be preceded by X or B and be enclosed in single quotation marks. The maximum number that can be specified is 9999; zero indicates that all entries are to be listed.

Abbreviation: CREAT

**ENTRIES(entryname[/password][b entryname[/password]...])|
LEVEL(level)**

specifies the names of entries to be listed. The entire catalog is listed when all of the following are true:

- ENTRIES, LEVEL, NOTUSABLE, CREATION, and EXPIRATION are not coded.
- No entrytype parameter is coded.
- TSO is not being used to issue the LISTCAT command.

Entries are only listed from an OS CVOL if the ENTRIES or LEVELS parameter is specified and:

- The CATALOG parameter is not specified
- No JOB or STEPCAT DD is specified
- The higher level qualifier of the name specified using ENTRIES or LEVELS is the alias of a CVOL

**ENTRIES(entryname[/password]
[b entryname[/password]...])**

specifies the name or generic name of each entry to be listed. When you want to list the entries that describe a user catalog, the catalog's volume must be physically mounted. You then specify the catalog's name as the entryname. If you want data space information, you must specify the volume serial number (as the entryname) of the volume containing the data space; you must also specify SPACE and no other entrytype parameters.

If the ENTRIES parameter specifies a generic name, a maximum of 1456 entries can be listed.

Note to TSO users: When you specify an incomplete qualified entryname, TSO prompts you to complete it.

password

specifies a password when the entry to be listed is password protected and a password is not specified with the CATALOG parameter. The password must be any

of the entry's passwords. The entry's protection attributes are listed only when you specify the entry's (or its catalog's) master password. When you don't supply a password for a password-protected entry, the operator or TSO terminal user is prompted for the entry's password. You cannot supply a password for these types of entries: nonVSAM data set, generation data group, alias, user-catalog connector, and data space.

Abbreviation: ENT

LEVEL(level)

specifies that all entries that match the level of qualification specified by (level) are to be listed irrespective of the number of qualifications after the level. If a generic level name is specified, one qualifier replaces the *.

Note: The * does not necessarily specify the last (rightmost) level of qualifications.

As an example, suppose the VSAM catalog contains the following names:

1. A.A.B
2. A.B.B
3. A.B.B.C
4. A.B.B.C.C
5. A.C.C
6. A.D
7. A.E

If ENTRIES(A.*) is specified, then entries 6 and 7 would be listed. If ENTRIES(A.*.B) is specified, then entries 1 and 2 would be listed. If LEVEL(A.*.B) is specified, then entries 1,2,3 and 4 would be listed. If LEVEL(A) is specified, then entries 1,2,3,4,5,6 and 7 would be listed.

If the LEVEL parameter is specified, a maximum of 1456 entries can be listed.

Note to TSO users: TSO will prefix the userid to the specified data-set name when the ENTRIES parameter is specified with an unqualified entry name. The userid is not prefixed when the LEVEL parameter is specified.

Abbreviation: LVL

EXPIRATION(days)

specifies that entries of the indicated type (CLUSTER, DATA, etc.) are to be listed only if they will expire the specified number of days from now or earlier.

days

specifies the number of days from now. This value can be expressed in decimal, hexadecimal, or binary. If it is expressed in hexadecimal or binary, it must be preceded by X or B and be enclosed in single quotation marks. The maximum number that can be specified is 9999 and it indicates that all entries are to be listed. Any value that exceeds the year 2000 will default to 99.999 (yyddd). Zero indicates that only entries that have already expired are to be listed.

Abbreviation: EXPIR

NOTUSABLE

specifies that only those data and index entries which have the "unusable" indicator on are to be listed. A data or index component is marked "unusable" when a system failure occurs that results in damage to the entry's cataloged information.

When the cataloged information is reset, the damaged entry and its backup copy (in the catalog recovery area) might not match when the space allocation information is compared. VSAM marks the catalog entry "unusable" until the space allocation information is corrected. See "Restoring Catalog Entries After System Failure" for more details.

Abbreviation: NUS

OUTFILE(dname)

specifies a data set, other than the SYSPRINT data set, to receive the output produced by LISTCAT (that is, the listed catalog entries). Completion messages produced by Access Method Services are sent, along with your job's JCL and input statements, to the SYSPRINT data set.

dname identifies a DD statement that describes the alternate output data set. If OUTFILE is not specified, the entries are listed in the SYSPRINT data set. If an alternate data set is specified, it must meet the requirements shown under "Output Data Sets" in the "Introduction."

Abbreviation: OFILE

LISTCAT EXAMPLES

Listing a Key-Sequenced Cluster's Entry: Example 1

In this example, a key-sequenced cluster entry is listed.

```
//LISTCAT1 JOB      ...
//STEP1   EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD       *
          LISTCAT -
          ENTRIES(D40.EXAMPLE.KSDS1) -
          CLUSTER -
          ALL
/*
```

The job control language statements are:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The LISTCAT output will also be sent to the SYSPRINT output device.

The LISTCAT command lists the cluster's catalog entry. It is assumed that the high level of the qualified cluster name is the same as the alias of the catalog D27UCAT1; this naming convention directs the catalog search to the appropriate catalog. Its parameters are:

- ENTRIES, which identifies the entry to be listed.
- CLUSTER, which specifies that only the cluster entry is to be listed. If CLUSTER had not been specified, the cluster's data and index entries would also be listed.
- ALL, which specifies that all fields of the cluster entry are to be listed.

Alter a Catalog Entry, Then List the Modified Entry: Example 2

In this example, the freespace attributes for the data component (KSDATA) of cluster MYDATA are modified. Next, the cluster entry, data entry, and index entry of MYDATA are listed to determine the effect, if any, the modification has on the cluster's other attributes and specifications.

```
//LISTCAT2 JOB      ...
//JOB CAT  DD       DSNAME=D27UCAT1,DISP=SHR
//STEP1   EXEC     PGM=IDCAMS
//SYSPRINT DD     SYSOUT=A
//SYSIN   DD       *
          ALTER -
          KSDATA -
          FREESPACE(10 10)
          IF LASTCC = 0 -
          THEN -
          LISTCAT -
          ENTRIES(MYDATA) -
          ALL
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The LISTCAT output will also be sent to the SYSPRINT output device.

The ALTER command modifies the freespace specifications of the key-sequenced VSAM cluster MYDATA. The command's parameters are:

- KSDATA, which is the entryname of the data component being altered. KSDATA identifies the data component of a key-sequenced VSAM cluster, MYDATA. In order to alter a value that applies only to the cluster's data component, such as FREESPACE does, you must specify the data-component's entryname.
- FREESPACE, which specifies the new freespace percentages for the data-component's control intervals and control areas.

The IF ... THEN command sequence verifies that the ALTER command completed successfully before the LISTCAT command executes.

The LISTCAT command lists the cluster's entry and its data and index entries. Its parameters are:

- ENTRIES, which specifies the entryname of the object being listed. Because MYDATA is a key-sequenced cluster, the cluster entry, its data entry, and its index entry are listed.
- ALL, which specifies that all fields of each entry are to be listed.

List Catalog Entries: Example 3

In this example, each catalog entry with the name "GENERIC.*.ABLE" is listed, where "*" is any 1 to 8 character simple name. The name "GENERIC.*.ABLE" is a generic name, and this example illustrates how all catalog entries with the same generic name are listed.

```
//LISTCAT3 JOB ...
//JOB CAT DD DSNAME=D27UCAT1,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
LISTCAT -
ENTRIES(GENERIC.*.ABLE) -
ALL
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT1.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer and the LISTCAT output are to be sent.

The LISTCAT command lists each catalog entry with the name GENERIC.*.ABLE, where "*" is any 1 to 8 character simple name. Its parameters are:

- ENTRIES, which specifies the entryname of the object to be listed. Because GENERIC.*.ABLE is a generic name, more than one entry might be listed.
- ALL, which specifies that all fields of each entry are to be listed.

LISTCRA

The format of the LISTCRA command is:

LISTCRA	INFILE(dname[b dname...]) [CATALOG(catname[/password]b dname)] [COMPARE NOCOMPARE] [DUMP NAME SEQUENTIALDUMP] [MASTERPW(password)] [OUTFILE(dname)]
----------------	--

LISTCRA may be abbreviated: LISTR.

LISTCRA PARAMETERS

Required Parameters

INFILE(dname[b dname...])
identifies the DD statement(s) that describes the catalog recovery area's volume. You can list or compare more than one catalog recovery area. However, if COMPARE is specified, all volumes whose catalog recovery areas are listed must be owned by the same VSAM catalog.

Note: The dname list for the INFILE parameter is limited to 9 dnames.

Abbreviation: IFILE

Optional Parameters

CATALOG(catname[/password]b dname)
identifies the catalog that owns the volume(s) identified with the INFILE parameter. The CATALOG parameter is required when you specify the COMPARE option either with the DUMP or NAME output format. CATALOG cannot be specified with SEQUENTIALDUMP.

password

When the catalog is password protected, you must supply the catalog's master password. The password is required even though the master catalog's password may have been supplied in the MASTERPW parameter.

dname

When CATALOG is specified, you must describe and allocate the catalog to be compared with a DD statement. dname identifies the catalog's DD statement, and is required when you specify the CATALOG parameter. You cannot specify JOBCAT or STEPCAT as the dname. However, if the catalog being compared is not the VSAM master catalog, you must also identify the catalog by a STEPCAT or JOBCAT DD statement in addition to the DD statement referenced by dname.

If the catalog information is entered incorrectly, the NOCOMPARE default is taken, and all of the CRA records are listed.

Abbreviation: CAT

COMPARE|NOCOMPARE

specifies whether the list is to be limited to those entries in the catalog recovery area which do not match their corresponding catalog entries.

When you specify COMPARE, only those entries which do not match are listed. All of the catalog's entries are compared, except the catalog's self-describing records. You must identify the catalog (that contains the to-be-compared catalog entries) with the CATALOG parameter.

When you specify NOCOMPARE (or allow it to default), all of the catalog recovery area's records are listed.

The COMPARE parameter cannot be specified if you specify SEQUENTIALDUMP.

Note: If you specify the COMPARE option, page space entries will always result in a mismatch. This is a normal condition and does not require any corrective action. The mismatch is due to the OPEN indicator not being set in the page space record in the CRA, whereas it is set in the catalog.

See "Appendix D: Interpreting LISTCRA Output Listings" for examples of each type of output listing.

Abbreviations: CMPR and NCMPR

DUMP|NAME|SEQUENTIALDUMP

specifies the amount of cataloged information to be listed.

DUMP

specifies that each listed entry is printed in its entirety in both hexadecimal and character form, sorted alphanumerically and grouped.

NAME

specifies that each listed entry includes only the entry's name, its volume serial numbers, and the name and entry type of each associated entry. All listed entries are sorted alphanumerically and grouped.

SEQUENTIALDUMP

specifies that each listed entry is to be printed in its entirety in both hexadecimal and character form. Records are listed in the sequence each appears in the catalog recovery area.

Abbreviation: SDUMP

The DUMP, NAME, SEQUENTIALDUMP, COMPARE, and NOCOMPARE options can be specified to produce five different kinds of output listing. See "Appendix D: Interpreting LISTCRA Output Listings" for examples of each type of listing.

MASTERPW(password)

specifies the master catalog's master password. You must supply the master catalog's master password with the MASTERPW parameter when the master catalog is password-protected.

Abbreviation: MRPW

OUTFILE(dname)

identifies the DD statement that describes an alternate output data set. See "Output Data Sets" for more details about alternate output data sets. When OUTFILE is not specified, the listing is printed on the output device described with the SYSPRINT DD statement. When OUTFILE is specified, the listing produced by LISTCRA is sent to the alternate output data set; the remainder of the Access Method Services output is printed on the output device described with the SYSPRINT DD statement. Refer to "Appendix D: Interpreting LISTCRA Output Listings" for an example of the use of the OUTFILE parameter.

Abbreviation: OFILE

LISTCRA EXAMPLE

Listing a Catalog Recovery Area: Example 1

This example lists the catalog recovery areas for those volumes owned by the VSAM master catalog, AMASTCAT, in dump format, and compares those catalog recovery areas with the actual catalog records themselves.

```
//LISTAREA JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CRAVOL1 DD DISP=OLD,VOL=SER=SG2001,UNIT=2314
//CRAVOL2 DD DISP=OLD,VOL=SER=VSER04,UNIT=3330
//CATVOL DD DSNAME=AMASTCAT,DISP=OLD
//SYSIN DD *
LISTCRA -
INFIL( -
CRAVOL1 -
CRAVOL2) -
MASTERPW(MCATMRPW) -
COMPARE -
DUMP -
CATALOG(AMASTCAT/MCATMRPW CATVOL)
/*
```

The job control language statements are:

- CRAVOL1 DD, which identifies and allocates the volume that contains the first catalog recovery area to be listed.
- CRAVOL2 DD, which identifies and allocates the volume that contains the second catalog recovery area to be listed.
- CATVOL DD, which identifies and allocates the catalog to be compared against, and makes the catalog available to LISTCRA as a data set. Because this example is listing the master catalog, no STEPCAT or JOBCAT DD statement is required. However, if the example were listing a user catalog, a STEPCAT or JOBCAT DD statement would be required in addition to the CATVOL DD statement to describe the user catalog as a catalog.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The listing produced by LISTCRA is also sent to this same output device.

The LISTCRA command causes the catalog recovery areas on volumes SG2001 and VSER04 to be listed and their contents compared with the VSAM system catalog AMASTCAT. The LISTCRA command's parameters are:

- INFIL, which is required and specifies the catalog recovery areas to be listed by identifying the DD statement that describes each CRA's volume.
- MASTERPW, which specifies the master password of the master catalog. This parameter is required in order to open the catalog recovery areas for LISTCRA processing.
- COMPARE, which specifies that the catalog records are to be compared with their copies in the catalog recovery areas.
- DUMP, which specifies that the results of the record comparison (that is, the LISTCRA output listing) is to be printed in dump format.
- CATALOG, which is required (because the COMPARE option was specified) and identifies the catalog to be compared against by identifying the DD statement that describes and allocates the catalog as a data set.

PRINT

The format of the PRINT command is:

PRINT	<pre>{INFILE(dname[/password]) INDATASET(entryname[/password])} [OUTFILE(dname)] [CHARACTER DUMP HEX] [FROMKEY(key) FROMADDRESS(address) FROMNUMBER(number) SKIP(count)] [TOKEY(key) TOADDRESS(address) TONUMBER(number) COUNT(count)]</pre>
--------------	---

PRINT PARAMETERS**Required Parameters**

**INFILE(dname[/password])|
INDATASET(entryname[/password])**
identifies the data set or component to be printed.

INFILE(dname[/password])
specifies the name of the DD statement that identifies the data set or component. You can list a base cluster in alternate-key sequence by specifying a path name as the data set name in the DD statement. See the section "Printing Data Sets" for information on DCB parameters required for nonVSAM data sets.

Abbreviation: IFILE

INDATASET(entryname[/password])
specifies the name of the entry to be printed. If INDATASET is specified, the entryname is dynamically allocated. You can list a base cluster in alternate-key sequence by specifying a path name as entryname.

password
If a VSAM data set or component is password protected, a password must be supplied. The password to be supplied is the master password of the catalog if you are listing a catalog, or the read- or higher-level password of the data set or component if the data set or component is not a catalog. You can specify the master password of the cluster if you are listing a component of a password-protected cluster. Passwords are applicable only to VSAM data sets and their components.

Abbreviation: IDS

Optional Parameters

CHARACTER|DUMP|HEX
specifies the format of the listing.

CHARACTER
specifies that each byte in the logical record is to be printed as a character. Bit patterns not defining a character are printed as periods. Key fields are listed in character format (see Figure 10).

Abbreviation: CHAR

listing begins with the first logical record in the data set or component. The only value that can be specified for a SAM data set is SKIP.

FROMKEY(key)

specifies the key of the first record you want listed. You can specify generic keys—that is, keys shorter than that defined for the data set. If you specify generic keys, listing begins at the first record whose key matches that portion of the key you specified. (You cannot specify a key longer than that defined for the data set. If you do, the listing is not performed.) If the specified key is not found, the next higher key is used as the starting point for the listing.

FROMKEY can be specified only when an alternate index, a key-sequenced VSAM data set, a VSAM catalog, or an indexed-sequential (ISAM) nonVSAM data set is being printed. When you specify FROMKEY, you cannot also specify TOADDRESS to identify the last record to be printed.

key can contain 1 to 255 EBCDIC characters. key must be enclosed in single quotation marks if it contains commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in single quotation marks. You can code key in hexadecimal form where the hexadecimal value is enclosed in single quotation marks and preceded by an X—for example, X'C2C1C8'. key can contain up to 255 hexadecimal characters.

Abbreviation: FKEY

FROMADDRESS(address)

specifies the relative byte address (RBA) of the first record you want listed. The RBA value must be the beginning of a logical record. If you specify this parameter for a key-sequenced data set, the listing will be in physical sequential order instead of in logical sequential order. FROMADDRESS can be specified only for VSAM key-sequenced or entry-sequenced data sets or components. FROMADDRESS cannot be specified when the data set is accessed through a path. FROMADDRESS cannot be specified for a key-sequenced data set with spanned records if any of those spanned records are to be accessed.

Abbreviation: FADDR

FROMNUMBER(number)

specifies the relative-record number of the first record you want printed.

Abbreviation: FNUM

SKIP(count)

specifies the number of logical records you want to skip before the listing of records begins. For example, if you want the listing to begin with record number 500, you specify SKIP(499). SKIP should not be specified when you are accessing the data set through a path; the results are unpredictable.

address, number, and count can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n'); the expression cannot be longer than one fullword (eight decimal or hexadecimal numbers, or 32 binary numbers).

OUTFILE(dname)

identifies an alternate output data set—that is, an output data set other than SYSPRINT. For dname substitute the name

of the JCL statement that identifies the alternate output data set. The standard Access Method Services output data set for listings, which is identified by the DD name SYSPRINT, is the default. The output data set must meet the requirements stated in the "Introduction" under "Output Data Sets."

Abbreviation: OFILE

**TOKEY(key)|TOADDRESS(address)|
TONUMBER(number)|COUNT(count)**

specifies the location in the data set being listed at which listing is to stop. If no value is specified, the listing ends with the logical end of the data set or component. The only value that can be specified for a sequential data set is COUNT. The location at which the listing is to stop must follow the location at which the listing is to begin.

The ending delimiter must be consistent with the starting delimiter. For example, if FROMADDRESS is specified for the starting location, neither TOKEY nor TONUMBER may be specified for the location. Similarly, if FROMNUMBER is specified for the starting location, neither TOKEY nor TOADDRESS may be specified for the ending location.

TOKEY(key)

specifies the key of the last record to be listed. You can specify generic keys—that is, keys shorter than that defined for the data set. If you specify generic keys, listing stops after the last record is listed whose key matches that portion of the key you specified. (You cannot specify a key longer than that defined for the data set. If you do, the listing is not performed.) If the specified key is not found, the next lower key is used as the stopping point for the listing.

TOKEY can be specified only when an alternate index, a key-sequenced VSAM data set, a VSAM catalog, or an indexed-sequential (ISAM) nonVSAM data set is being printed.

key can contain 1 to 255 EBCDIC characters. key must be enclosed in single-quotation marks if it contains commas, semicolons, blanks, parentheses, or slashes. A single-quotation mark within a key must be coded as two single-quotation marks if the key is enclosed in single quotation marks. You can code key in hexadecimal form where the hexadecimal value is enclosed in single-quotation marks and preceded by an X—for example, X'C2C1C8'. key can contain up to 255 hexadecimal characters.

TOADDRESS(address)

specifies the relative byte address (RBA) of the last record you want listed. Unlike FROMADDRESS, the RBA value does not need to be the beginning of a logical record. The entire record containing the specified RBA is printed. If you specify this parameter for a key-sequenced data set, the listing will be in physical sequential order instead of in logical sequential order. TOADDRESS can be specified only for VSAM key-sequenced or entry-sequenced data sets or components. TOADDRESS cannot be specified when the data set is accessed through a path. TOADDRESS cannot be specified for a key-sequenced data set with spanned records if any of those spanned records are to be accessed.

Abbreviation: TADDR

TONUMBER(number)
specifies the relative-record number of the last record you want printed.

Abbreviation: TNUM

COUNT(count)
specifies the number of logical records to be listed. COUNT should not be specified when you are accessing the data set through a path; the results are unpredictable.

address, number, and count can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n'); the expression cannot be longer than one fullword (eight decimal or hexadecimal numbers, or 32 binary numbers).

PRINT EXAMPLES

Print a Key-Sequenced Cluster's Data Records: Example 1

In this example, the data records of a key-sequenced cluster, D40.EXAMPLE.KSDS1, are printed in dump format. That is, each character of the record is printed in its hexadecimal and alphameric forms.

```
//PRINT1 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
PRINT -
          INDATASET(D40.EXAMPLE.KSDS1)
/*
```

The job control language statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The listing produced by PRINT is sent to the same output device.

The PRINT command prints data records of the key-sequenced cluster, D40.EXAMPLE.KSDS1. Its parameter is:

- INDATASET, which names the data set to be printed. Since the data set is cataloged in D27UCAT1 and the high-level qualifier of the data set is the name of the alias of D27UCAT1, Access Method Services can dynamically allocate the cluster without the need of JOBCAT or STEPCAT DD statements.

Because neither FROMADDRESS, FROMKEY, SKIP, TOKEY, TOADDRESS, or COUNT is specified, Access Method Services assumes that all of the cluster's data records are to be printed.

Because neither HEX nor CHAR was specified, Access Method Services prints each record in the DUMP format. An example of the printed record is shown in Figure 13.

```
KEY OF RECORD - 00F0F0F0F0F1C9E240C4C1405CC6C9
0000 00F0F0F0 F0F1C9E2 40C4C140 5CC6C9D3 C540C9F0 C6F8F05C 40F5F040 D9C5C3D6 *.00001IS DA *FILE I0D80* 50 RECO*
0020 D9C4E240 D6C640F6 F940C3C8 C1D9E240 E6C9E3C8 40D2C5E8 40C9D540 D7D6E240 *RDS OF 69 CHARS WITH KEY IN POS *
0040 F160F1F1 4B000000 00000000 00000000 *1-11.....
```

Figure 13. An Example of the Printed Record in DUMP Format

Copy Records From a NonVSAM Data Set Into an Entry-Sequenced VSAM Cluster, Then Print the Records: Example 2

In this example, the first fifteen records from a nonVSAM data set, EXAMPLE.NONVSAM, are copied into an entry-sequenced cluster, D50.EXAMPLE.ESDS1. If the records were copied correctly, the cluster's records are printed in hexadecimal format. Finally, even if the records were not copied correctly, the nonVSAM data set's first fifteen records are printed in character format.

```
//PRINT2 JOB ...
//JOB CAT DD DSNAME=D27UCAT1,DISP=SHR
// DD DSNAME=D27UCAT2,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//VSDSET2 DD DSNAME=D50.EXAMPLE.ESDS1,DISP=OLD
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
      REPRO -
          INDATASET(EXAMPLE.NONVSAM) -
          OUTFILE(VSDSET2/DEPT27U) -
          COUNT(15)

      IF LASTCC = 0 -
          THEN -
          PRINT -
              INFILE(VSDSET2) -
              HEX

      PRINT -
          INDATASET(EXAMPLE.NONVSAM) -
          COUNT(15) -
          CHARACTER
/*
```

The job control language statements are:

- JOBCAT DD, which makes two catalogs available for this job: D27UCAT1 and D27UCAT2. Concatenated JOBCAT DD statements were used to identify both catalogs.
- VSDSET2 DD, which identifies the entry-sequenced VSAM cluster, D50.EXAMPLE.ESDS1, that the records are copied into.

Note: If the AMP=(BUFND=n) parameter was specified, performance would improve when the data set's records are accessed. BUFND was allowed to default in this example, because only 15 records are being processed.

- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The listing produced by the PRINT commands is sent to the same output device.

The REPRO command copies the first fifteen records from the input data set, EXAMPLE.NONVSAM, into the output entry-sequenced cluster, D50.EXAMPLE.ESDS1. Its parameters are:

- INDATASET, which identifies the input data set, EXAMPLE.NONVSAM. Because a JOBCAT DD statement is included with the job, VSAM will search for the catalog entry describing the nonVSAM data set in either D27UCAT1 or D27UCAT2 user catalog, or in the master catalog.
- OUTFILE, which points to the VSDSET2 DD statement. The VSDSET2 DD statement identifies the output data set, D50.EXAMPLE.ESDS1. VSAM assumes that D50.EXAMPLE.ESDS1 is cataloged either in D27UCAT1 or D27UCAT2 user catalog, or in the master catalog.

Print a VSAM Catalog: Example 3

This example shows various ways to print a VSAM catalog. You may find this function of the PRINT command helpful in the event of a problem with your catalog.

```
//PRINT3   JOB   ...
//JOB CAT  DD   DSNAME=D27UCAT1,DISP=OLD
//STEP1    EXEC PGM=IDCAMS
//SYS PRINT DD  SYSOUT=A
//SYS IN   DD   *

/* PRINT THE LOW KEY RANGE OF THE CATALOG */
PRINT -
  INDATASET(D27UCAT1) -
  TOKEY (X'3F')

/* PRINT THE HIGH KEY RANGE OF THE CATALOG */
PRINT -
  INDATASET(D27UCAT1) -
  FROMKEY (X'40')

/* PRINT A SPECIFIC RECORD */
/* FROM THE LOW KEY RANGE */
PRINT -
  INDATASET(D27UCAT1) -
  FROMKEY (X'0000000E') -
  COUNT (1)

/* PRINT THE ENTIRE CATALOG */
PRINT -
  INDATASET(D27UCAT1)
/*
```

The job control language statements are:

- JOBCAT DD, which is required and which describes and allocates the user catalog to be printed.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent. The listings produced by the PRINT commands are sent to the same output device.

The first PRINT command prints the low-keyrange of the catalog. Its parameters are:

- INDATASET, which specifies the name of the catalog to be printed.
- TOKEY, which specifies that printing is to terminate after reaching a record whose key is greater than X'3F' (in its first byte). The low-keyrange of a VSAM catalog contains records with a range of keys having a value from X'00' to X'3F' in the first byte.

The second PRINT command prints the high-keyrange of the catalog. Its parameters are:

- INDATASET, which specifies the name of the catalog to be printed.
- FROMKEY, which specifies that printing is to begin at the first record whose key is X'40' or greater (in its first byte). The high-keyrange of a VSAM catalog contains records with a range of keys having a value from X'40' to X'FF' in the first byte.

The third PRINT command prints one catalog record from the low-keyrange of the catalog. Its parameters are:

- INDATASET, which specifies the name of the catalog to be printed.
- FROMKEY, which specifies that printing is to begin with the record whose key is X'0000000E' in the first four bytes. Records in the low-keyrange of a VSAM catalog have the number of the containing control interval as the first four bytes of their key. Records are 505 bytes in length and each record is contained in a 512-byte control interval.
- COUNT, which specifies that only one record is to be printed.

The fourth PRINT command prints the entire catalog, because there are no delimiting parameters specified.

REPRO

The format of the REPRO command is:

REPRO	<pre>{INFILE(dname[/password] [b ENVIRONMENT(DUMMY)]) INDATASET(entryname[/password] [b ENVIRONMENT(DUMMY)])} {OUTFILE(dname[/password]) OUTDATASET(entryname[/password])} [REPLACE NOREPLACE] [REUSE NOREUSE] [FROMKEY(key) FROMADDRESS(address) FROMNUMBER(number) SKIP(count)] [TOKEY(key) TOADDRESS(address) TONUMBER(number) COUNT(count)]</pre>
--------------	---

REPRO PARAMETERS**Required Parameters**

INDATASET(entryname[/password][b ENVIRONMENT(DUMMY)])
 identifies the data set to be copied. See the section "Copying Data Sets" for information on the types of data sets which can be copied, specific copy functions, and restrictions. See also the section "Copying a Catalog."

INFILE(dname[/password])
 specifies the name of the DD statement that identifies the data set to be copied. You can copy a base cluster in alternate-key sequence by specifying a path name as the data set name in the DD statement. See the section "Copying Data Sets" for information on DCB parameters required for nonVSAM data sets.

Abbreviation: IFILE

INDATASET(entryname[/password])
 specifies the name of the entry to be copied. If INDATASET is specified, the entryname is dynamically allocated. You can copy a base cluster in alternate-key sequence by specifying a path name for entryname.

password
 is the read- or higher-level password of the data to be copied. If the data set is password protected, the read password must be supplied. If a catalog is to be copied, the master password is required. Passwords are applicable only to VSAM data sets.

Abbreviation: IDS

ENVIRONMENT (DUMMY)
 specifies that dummy ISAM records are to be copied. Dummy records are records with hexadecimal 'FF' in the first byte. If you do not include this parameter, dummy records will be ignored during the copy operation. See OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide for further information. If ENVIRONMENT is not coded, dummy records are not copied.

Abbreviations: ENV and DUM

OUTFILE(dname[/password])|

OUTDATASET(entryname[/password])
identifies the output data set. ISAM data sets cannot be specified as output data sets.

OUTFILE(dname[/password])
specifies the name of a DD statement that identifies the output data set. For VSAM data sets, the data set name can be that of a path. See the section "Copying Data Sets" for information on DCB parameters required for nonVSAM data sets. If the DD statement identifies a SYSOUT data set, the attributes must match those specified in the section: "Output Data Sets."

Abbreviation: OFILE

OUTDATASET(entryname[/password])
specifies the name of the output data set. If OUTDATASET is specified, the entryname is dynamically allocated. For VSAM data sets, entryname can be that of a path.

Abbreviation: ODS

password
specifies the update- or higher-level password for a password-protected output data set or path.

Optional Parameters

**FROMKEY(key)|FROMADDRESS(address)|
FROMNUMBER(number)|SKIP(count)**
specifies the location in the input data set from which copying is to start. You can use only one of the four possible choices. If no value is coded, the listing begins with the first logical record in the data set. The only parameter that can be coded for a SAM data set is SKIP. None of these parameters can be specified if you are copying a catalog; the entire catalog must be copied.

FROMKEY(key)
specifies the key of the first record you want copied. You can specify generic keys—that is, keys shorter than that defined for the data set. If you specify generic keys, copying begins at the first record whose key matches that portion of the key you specified. (You cannot specify a key longer than that defined for the data set. If you do, the data set is not copied.) If the specified key is not found, the next higher key is used as the starting point for copying.

FROMKEY can be specified only when an alternate index, a key-sequenced VSAM data set, or an indexed-sequential (ISAM) nonVSAM data set is being copied. If FROMKEY is specified for the starting location, TOADDRESS cannot identify the ending location.

key can contain 1 to 255 EBCDIC characters. key must be enclosed in single quotation marks if it contains commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in single quotation marks. You can code key in hexadecimal form where the hexadecimal value is enclosed in single quotation marks and preceded by an X—for example, X'C2C1C8'. key can contain up to 255 hexadecimal characters.

Abbreviation: FKEY

FROMADDRESS(address)

specifies the relative byte address (RBA) of the first record you want copied. The RBA value must be the beginning of a logical record. If you specify this parameter for key-sequenced data, the records will be copied in physical sequential order instead of in logical sequential order. FROMADDRESS can be coded only for VSAM key-sequenced or entry-sequenced data sets or components. FROMADDRESS cannot be specified when the data set is being accessed through a path. FROMADDRESS cannot be specified for a key-sequenced data set with spanned records if any of those spanned records are to be accessed.

Abbreviation: FADDR

FROMNUMBER(number)

specifies, for a relative-record cluster, the relative-record number of the first record you want copied.

Abbreviation: FNUM

SKIP(count)

specifies the number of logical records you want to skip before beginning to copy records. For example, if you want to copy beginning with record number 500, you specify SKIP(499). SKIP should not be specified when you access the data set through a path; the results are unpredictable.

address, number, and count can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form; the expression cannot be longer than one fullword (eight decimal or hexadecimal numbers, or 32 binary numbers).

REPLACE|NOREPLACE

specifies whether a record in the source cluster (INFILE or INDATASET) is to replace a record in the target cluster (OUTFILE or OUTDATASET) when the source cluster is copied into the target cluster. When a key-sequenced cluster or relative-record cluster (that is, the source cluster) is copied, its records might have keys or relative-record numbers identical to the keys or relative-record numbers of data records in the target cluster.

REPLACE cannot be used if the output data set is identified as a path through an alternate index, or if the output data set is a base cluster whose upgrade set includes an alternate index defined with the unique-key attribute.

Since the catalog Reload function automatically performs a Replace function, REPLACE|NOREPLACE is ignored if the output data set is a VSAM catalog.

When a key-sequenced cluster, other than a catalog, is copied and you specify:

- REPLACE, each source record whose key matches a target record's key replaces the target record. Otherwise, the source record is inserted into its appropriate place in the target cluster.
- NOREPLACE, target records aren't replaced by source records. Each source record whose key matches a target record's key is not inserted into the target data set and a "duplicate record" message is issued.

When a relative-record cluster is copied and you specify:

- REPLACE, each source record whose relative-record number identifies a data record (rather than an empty slot) in the target cluster replaces the target data record. Otherwise, the source data record is inserted into the empty slot its relative-record number identifies.
- NOREPLACE, target records aren't replaced by source records. Each source record whose relative-record number identifies a target data record instead of an empty slot is not inserted into the target data set and a "duplicate record" message is issued.

Abbreviations: REP and NREP

REUSE|NOREUSE

specifies whether the output data set is to be opened as a reusable data set. This parameter is valid only for VSAM data sets. When REUSE is specified, the output data set, specified with OUTFILE or OUTDATASET, is opened as a reusable data set regardless of whether or not it was defined as REUSE (see the DEFINE CLUSTER command description). If the data set was defined as REUSE, its high-used relative byte address (RBA) is reset to zero (that is, the data set is effectively empty) and the operation proceeds. If REUSE is specified and the data set was originally defined with the NOREUSE option, the data set must be empty; otherwise, the REPRO command terminates with an error message.

If the NOREUSE attribute is specified (either explicitly or by default) and OUTFILE or OUTDATASET identifies a nonempty data set, records are written at the end of an entry-sequenced data set. Processing of key-sequenced and relative-record data sets is controlled by the REPLACE/NOREPLACE option used.

Abbreviations: RUS and NRUS

**TOKEY(key)|TOADDRESS(address)|
TONUMBER(number)|COUNT(count)**

specifies the location in the data set being copied at which copying is to stop. You can use only one of the four choices. If no value is coded, the copying ends with the logical end of the data set or component. The only parameter that can be specified for a sequential data set is COUNT. None of these parameters can be specified if you are copying a catalog; the entire catalog must be copied. The location at which the copying is to end must follow the location at which it is to begin.

TOKEY(key)

specifies the key of the last record you want copied. You can specify generic keys—that is, keys shorter than that defined for the data set. If you specify generic keys, copying stops after the last record is copied whose key matches that portion of the key you specified. (You cannot specify a key longer than that defined for the data set. If you do, the data set is not copied.) If the specified key is not found, the next lower key is used as the stopping point for copying.

TOKEY can be specified only when an alternate index, a key-sequenced VSAM data set, or an indexed-sequential (ISAM) nonVSAM data set is being copied. If FROMADDRESS was specified for the starting location, TOKEY cannot be specified for the ending location.

key can contain 1 to 255 EBCDIC characters. 'key must be enclosed in single quotation marks if it contains commas, semicolons, blanks, parentheses, or slashes. A single quotation mark within a key must be coded as two single quotation marks if the key is enclosed in

single quotation marks. You can code key in hexadecimal form where the hexadecimal value is enclosed in single quotation marks and preceded by an X—for example, X'C2C1C8'. key can contain up to 255 hexadecimal characters.

TOADDRESS(address)

specifies the relative byte address (RBA) of the last record you want copied. Unlike FROMADDRESS, the RBA value does not need to be the beginning of a logical record. The entire record containing the specified RBA is copied. If you specify this parameter for a key-sequenced data set, the records will be copied in physical sequential order instead of in logical sequential order.

TOADDRESS cannot be specified as the ending location if FROMKEY was specified as the starting address. TOADDRESS can be used only with VSAM key-sequenced or entry-sequenced data sets or components. TOADDRESS cannot be specified when the data set is accessed through a path. TOADDRESS cannot be specified for a key-sequenced data set with spanned records if any of those spanned records are to be accessed.

Abbreviation: TADDR

TONUMBER(number)

specifies the relative-record number of the last record you want copied. TONUMBER can be specified only when you copy a relative-record cluster.

Abbreviation: TNUM

COUNT(count)

specifies the number of logical records you want copied. COUNT should not be specified when you access the data set through a path; the results are unpredictable.

address, number, and count can be expressed in decimal (n), hexadecimal (X'n'), or binary (B'n') form; the expression cannot be longer than one fullword (eight decimal or hexadecimal numbers, or 32 binary numbers).

REPRO EXAMPLES

Copy Records Into a VSAM Cluster: Example 1

In this example, records from an indexed-sequential data set, ISAMDSET, are copied into key-sequenced VSAM cluster, D40.EXAMPLE.KSDS1.

```
//REPRO1 JOB ...
//JOB CAT DD DSN= D27UCAT2,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//INDSET1 DD DSN=ISAMDSET,DISP=OLD,
// DCB=(DSORG=IS,BUFNO=6)
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
REPRO -
INFIL(INDSET1) -
OUTDATASET(D40.EXAMPLE.KSDS1)
/*
```

The job control language statements are:

- JOBCAT DD, which makes a catalog available for this job: D27UCAT2.
- INDSET1 DD, which describes the indexed-sequential data set, ISAMDSET, that contains the input records. The BUFNO parameter specifies the number of buffers assigned to the ISAM data set. This improves performance when the ISAM data set's records are accessed.
- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The REPRO command copies all records from the input data set, ISAMDSET, to the output data set, D40.EXAMPLE.KSDS1. Its parameters are:

- INFIL, which points to the INDSET1 DD statement. The INDSET1 DD statement identifies the source, or input, data set: ISAMDSET. ISAMDSET is an indexed-sequential data set.
- OUTDATASET, which identifies the key-sequenced VSAM cluster into which the input records are to be copied. This data set is dynamically allocated by Access Method Services.

Copy Records Into a VSAM Cluster: Example 2

In this two-part example, data records are copied from the nonVSAM data set SEQ.D27V, a sequential data set, into a key-sequenced VSAM cluster, D40.MYDATA. Next, records are copied from the key-sequenced cluster, D40.MYDATA, into an entry-sequenced cluster, ENTRY.

```
//REPRO2 JOB ...
//JOB CAT DD DSN=MYCAT,DISP=SHR
//STEP1 EXEC PGM=IDCAMS
//INPUT DD DSN=SEQ.D27V,DISP=SHR,DCB=(BUFNO=6)
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
REPRO -
      INFILE(INPUT) -
      OUTDATASET(D40.MYDATA)
/*
//STEP2 EXEC PGM=IDCAMS
//INPUT DD DSN=D40.MYDATA,DISP=OLD
//OUTPUT DD DSN=ENTRY,DISP=OLD
//SYS PRINT DD SYSOUT=A
//SYS IN DD *
REPRO -
      INFILE(INPUT) -
      OUTFILE(OUTPUT) -
      FROMKEY(DEAN) -
      TOKEY(JOHNSON)
/*
```

STEP1: Access Method Services copies records from a sequential data set, SEQ.D27V, into a key-sequenced VSAM cluster, D40.MYDATA. STEP1's job control language statements are:

- JOBCAT DD, which makes the catalog MYCAT available for this job.
- INPUT DD, which identifies the sequential data set, SEQ.D27V, that contains the input records. The BUFNO parameter specifies the number of buffers assigned to the sequential data set. This improves performance when the data set's records are accessed.
- SYS PRINT DD, which is required in all Access Method Services job steps. The SYS PRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

STEP1's REPRO command copies all records from the input data set, SEQ.D27V, to the output data set, D40.MYDATA. Its parameters are:

- INFILE, which points to the INPUT DD statement. The INPUT DD statement identifies the source, or input, data set.
- OUTDATASET, which identifies the key-sequenced cluster into which the input records are to be copied. The data set is dynamically allocated by Access Method Services.

STEP2: Access Method Services copies some of the records of the key-sequenced cluster D40.MYDATA into another entry-sequenced cluster, ENTRY. STEP2's job control language statements are:

- JOBCAT DD, which applies to this job step as well as to STEP1.
- INPUT DD, which identifies the key-sequenced cluster, D40.MYDATA, that contains the input records.
- OUTPUT DD, which identifies the entry-sequenced cluster, ENTRY, that the records are copied into.

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

STEP2's REPRO command copies records from the input data set, D40.MYDATA, to the output data set, ENTRY. Only those records with key values from DEAN to, and including, JOHNSON are copied. The REPRO command's parameters are:

- INFILE, which points to the INPUT DD statement. The INPUT DD statement identifies the source, or input, key-sequenced cluster.
- OUTFILE, which points to the OUTPUT DD statement. The OUTPUT DD statement identifies the entry-sequenced cluster into which the input records are to be copied.
- FROMKEY and TOKEY, which specify the lower and upper key boundaries. Only those records with key values from DEAN to, and including, JOHNSON are copied.

If ENTRY already contains records, VSAM merges the copied records with ENTRY's records. A subsequent job step could resume copying the records into ENTRY, beginning with the records with key greater than JOHNSON. If you subsequently copied records with key values less than DEAN into ENTRY, VSAM merges them with ENTRY's records.

Copy a Catalog: Example 3

In this example, a catalog is copied to illustrate the copy-catalog procedure.

```

//COPYCAT JOB    ...
//STEP1  EXEC   PGM=IDCAMS
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          DEFINE USERCATALOG -
          (NAME(COPYUCAT) -
            MASTERPW(UCATMPW) -
            UDATEPW(UCATUPW) -
            FOR(365) -
            VOLUME(VSER06) -
            CYLINDERS(100 10) ) -
          DATA -
            (CYLINDERS(20 10) ) -
          INDEX -
            (CYLINDERS(10) ) -
          CATALOG -
            (AMASTCAT/MRCATPW2)
/*
//STEP2  EXEC   PGM=IDCAMS
//STEPCAT DD   DSNAME=COPYUCAT,DISP=OLD
//        DD   DSNAME=MYCAT,DISP=OLD
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          REPRO -
            INDATASET(MYCAT) -
            OUTDATASET(COPYUCAT/UCATMPW)
          EXPORT -
            MYCAT -
            DISCONNECT
/*
//STEP3  EXEC   PGM=IDCAMS
//STEPCAT DD   DSNAME=COPYUCAT,DISP=SHR
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          LISTCAT
/*
//STEP4  EXEC   PGM=IDCAMS
//STEPCAT DD   DSNAME=COPYUCAT,DISP=OLD
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          DELETE -
            MYCAT -
            CLUSTER -
            PURGE -
            CATALOG(COPYUCAT/UCATMPW)
/*
//STEP5  EXEC   PGM=IDCAMS
//SYSPRINT DD   SYSOUT=A
//SYSIN   DD    *
          DEFINE ALIAS -
            (NAME(MYCAT) -
            RELATE(COPYUCAT) )
/*

```

STEP1: A user catalog, COPYUCAT, is defined on volume VSER06. The catalog entries in MYCAT will be copied into COPYUCAT. STEP1's job control statement is:

- SYSPRINT DD, which is required in all Access Method Services job steps. The SYSPRINT DD statement identifies the output device to which Access Method Services messages to the programmer are to be sent.

The DEFINE USERCATALOG command defines the user catalog, COPYUCAT. Its parameters are:

- NAME, which specifies the name of the new catalog, COPYUCAT.

- MASTERPW and UPDATEPW, which specify the master and update level passwords for the catalog.
- FOR, which specifies that the catalog is to be retained for 365 days.
- VOLUME, which specifies that the catalog is to reside on volume VSER06.
- CYLINDERS, which specifies that the catalog's data space is initially to be 100 cylinders. When the data space is extended, it is to be extended in increments of 20 cylinders. Part of the data space contains the catalog, COPYUCAT. The rest of the data space is available for suballocated VSAM clusters.
- DATA(CYLINDERS) and INDEX(CYLINDERS) specify that the catalog itself is initially to occupy 30 cylinders. VSAM determines the proportion of space to allocate to the catalog's data and index components. When the catalog's data component is extended, it is to be extended in increments of 10 cylinders.
- CATALOG, which specifies that a user catalog connector entry is to be built and written into the AMASTCAT catalog. The user catalog connector entry points to COPYUCAT.

STEP2: Access Method Services copies the contents of MYCAT into COPYUCAT. STEP2's job control statements are:

- STEPCAT DD, which makes two catalogs available for this job step: COPYUCAT and MYCAT. The DD statements that identify the catalogs are concatenated.
- SYSPRINT DD (see STEP1 above).

The REPRO command copies all records from MYCAT into COPYUCAT. Access Method Services treats each catalog as a key-sequenced VSAM cluster and copies each catalog record. Consequently, the first thirteen catalog records of MYCAT, which describe MYCAT as a key-sequenced cluster, are also copied into COPYUCAT. Entries from MYCAT are written into COPYUCAT beginning with catalog record 14 (that is, after the thirteen self-describing records of COPYUCAT). The REPRO command's parameters are:

- INDATASET, which identifies the source, or input, data set: MYCAT. VSAM assumes that MYCAT is cataloged in either MYCAT or COPYUCAT.
- OUTDATASET, which identifies the receiving data set: COPYUCAT. VSAM assumes that COPYUCAT is cataloged in either COPYUCAT or MYCAT.

The EXPORT command removes MYCAT's user catalog connector entry from the master catalog. MYCAT's cataloged objects are now not available to the system. (STEP5 builds an alias entry that relates MYCAT to COPYUCAT, making the cataloged objects available to the system again.)

STEP3: Access Method Services lists the name of each entry in the new catalog, COPYUCAT. The STEPCAT DD statement identifies the catalog to be listed.

LISTCAT cannot run in a job step where the catalog is empty when it is opened. To ensure that the LISTCAT correctly reflects the contents of the catalog, the LISTCAT was run as a separate job step.

STEP4: Access Method Services removes the entries in COPYUCAT that describe MYCAT as a key-sequenced cluster. (See STEP2 description above.) The DELETE command's parameters are:

- MYCAT, which identifies the object to be deleted.

- CLUSTER, which specifies that MYCAT is cataloged as a cluster in COPYUCAT. MYCAT's cluster, data, and index entry are removed from COPYUCAT.
- PURGE, which specifies that MYCAT is to be deleted regardless of a specified retention period or date.
- CATALOG, which specifies that MYCAT is cataloged in COPYUCAT. COPYUCAT's master password, UCATMPW, is required to delete its catalog entries.

STEP5: Access Method Services builds an alias entry that relates MYCAT to COPYUCAT. Because no CATALOG parameter or JOBCAT or STEPCAT DD statement identifies the catalog that is to contain the alias entry, VSAM assumes the entry is to be written into the master catalog.

When MYCAT was defined, it was defined on volume VSER04. Its catalog entries describe VSAM objects on that volume. COPYUCAT is defined on volume VSER06. Because MYCAT was copied into COPYUCAT (and MYCAT no longer exists as a user catalog), COPYUCAT entries now describe VSAM objects on volumes VSER04 and VSER06.

Unload a VSAM User Catalog: Example 4

This example shows how a VSAM user catalog can be unloaded to tape using the catalog unload/reload feature of the REPRO command.

```
//UNLOAD JOB ...
//JOBCAT DD DSNAME=D27UCAT2,DISP=OLD
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CATIN DD DSNAME=D27UCAT2,DISP=OLD
//CATOUT DD DSNAME=PORTABLE.TAPE,LABEL=(1,SL),
// DISP=NEW,UNIT=2400-3,VOL=SER=TAPE01,
// DCB=(DEN=3,RECFM=VB,LRECL=516,BLKSIZE=5164)
//SYSIN DD *
REPRO -
      INFILE(CATIN/UCATMRPW) -
      OUTFILE(CATOUT)
/*
```

The job control language statements are:

- JOBCAT DD, which is required, and which identifies and allocates the user catalog as the job catalog.
- CATIN DD, which describes and allocates the user catalog, D27UCAT2, as a VSAM data set to be opened and used by the REPRO command as the source data set for the unload operation.
- CATOUT DD, which describes and allocates a magnetic tape file to contain the copy of the catalog. The DCB parameters must be specified as shown. The BLKSIZE parameter can be any multiple of LRECL, plus 4, where the minimum LRECL=505.

The REPRO command causes the VSAM user catalog D27UCAT2 to be unloaded (that is, copied) to a magnetic tape file. The REPRO command's parameters are:

- INFILE, which is required, and which identifies the user catalog as an input, or source, VSAM data set. The catalog's master password is required to open it as a data set.
- OUTFILE, which is required, and which describes the magnetic tape file, or target data set that is to contain the catalog's copy.

Reload an Unloaded VSAM User Catalog: Example 5

This example shows how a VSAM user catalog can be reloaded from the backup copy created in the previous example using the catalog unload/reload feature of the REPRO command.

```
//RELOAD JOB ...
//JOB CAT DD DSNAME=D27UCAT2,DISP=OLD
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CATIN DD DSNAME=PORTABLE.TAPE,LABEL=(1,SL),
// DISP=OLD,UNIT=2400-3,VOL=SER=TAPE01,DCB=DEN=3
//CATOUT DD DSNAME=D27UCAT2,DISP=OLD
//SYSIN DD *
REPRO -
      INFILE(CATIN) -
      OUTFILE(CATOUT/UCATMRPW)
/*
```

The job control language statements are:

- JOBCAT DD, which is required and which describes and allocates the user catalog to be reloaded. The JOBCAT statement specifies that D27UCAT2 is the job catalog.
- CATIN DD, which describes and allocates the magnetic tape file that contains the unloaded copy of the catalog. Since a standard label tape is used, DCB parameters of record format, recordsize, and blocksize need not be specified. For a non-labeled tape, the same parameters as in the previous example must be specified.
- CATOUT DD, which describes the user catalog as a VSAM data set to be opened and used by REPRO as the target data set for the reload operation.

The REPRO command causes a VSAM user catalog to be reloaded from the backup copy created in the previous example. The REPRO command's parameters are:

- INFILE, which is required, and which identifies the magnetic tape file that contains the unloaded, or backup, copy of the user catalog.
- OUTFILE, which is required and which identifies the user catalog as a VSAM data set to be opened for record processing. The master password allows opening the catalog as a data set.

RESETCAT

The format of the RESETCAT command is:

RESETCAT	CATALOG (catname[/password][b dname]) {CRAFILES ((dname1[b ALL b NONE]) [(dname1[b ALL b NONE])...]) CRAVOLUMES ((volser[b devtype]) [(volser[b devtype])...]) [IGNORE NOIGNORE] [MASTERPW (password) [WORKCAT (catname[/password]) [WORKFILE (dname[/password])
-----------------	---

RESETCAT may be abbreviated: RCAT.

RESETCAT PARAMETERS

Required Parameters

CATALOG(catname[/password][dname])
 identifies the catalog to be reset. The catalog specified by catname must have the RECOVERABLE attribute. Exclusive control of the catalog is required throughout the duration of the command; also, no VSAM data sets cataloged in the catalog may be open. The master catalog may not be reset while it is in use as a master catalog.

password
 specifies the catalog's master password. If the catalog is password-protected, you must specify the master password of the catalog.

dname
 specifies the name of the DD statement for the catalog being reset. If dname is not specified, the catalog will be dynamically allocated. It must also be specified in a JOBCAT or STEPCAT.

Abbreviation: CAT

CRAFILES((dname1[b ALL|b NONE])...)
 specifies a list of DD statements which identify the volumes and devices to be used to reset the catalog. Each volume contains a CRA which has a copy of the catalog records describing the VSAM objects on that volume. ALL specifies that the catalog records in the CRA on the volume specified by dname1 are to be reset. If a volume contains a multivolume data set whose catalog entry is being reset, other volumes of that multivolume data set may be needed during the reset operation; you should specify these via additional dname1 parameters. Use the NONE subparameter, if you do not want the entries in these additional volumes to be used to reset the catalog. If a volume is needed and you do not specify the DD statements provided by CRAFILES, it will be allocated dynamically, and will assume the NONE attribute.

CRAVOLUMES((volser[b devtype])...)
 specifies the volume serial numbers and corresponding device types of the volumes used in resetting the catalog. Each volume contains a catalog recovery area (CRA) which has a copy of the catalog records that describe the objects on the volume. The volume is allocated dynamically.

devtype
 specifies the device type for the associated volser. If you do not specify devtype, the volume must be mounted and on a unit marked permanently resident or

reserved. If you specify a generic name for devtype, the volume need not be mounted.

If a volume contains a multivolume data set whose catalog entry is being reset, other volumes of that multivolume data set may be needed during the reset operation; if needed and not specified for reset, these volumes will be allocated dynamically. Dynamically allocated volumes assume the NONE attribute as described for the CRAFILES parameter.

Abbreviation: CRAVOL

Optional Parameters

IGNORE|NOIGNORE

specifies whether RESETCAT should continue processing and force the reset when certain errors are encountered. These errors may be:

- I/O error in the catalog
- I/O error in the CRA
- The volume record or its extensions in the CRA are in error.

When you specify IGNORE, the RESETCAT command will try to recover as much information as possible and reset the catalog's description of those data sets on the specified volumes. For example, if IGNORE is specified and an I/O error occurs while processing a record from a CRA, that record will be ignored. In the process, data sets that could not be recovered on a particular volume would be marked unusable on that volume, and the space that was assigned to them would be freed. NOIGNORE specifies that these errors result in immediate termination.

Abbreviations: IGN and NIGN

MASTERPW(password)

specifies the master catalog's master password. The master catalog's master password is required when the master catalog is password protected.

WORKCAT(catname[/password])

specifies the name of the catalog in which RESETCAT defines the WORKFILE. This catalog must not be the same as the catalog being reset. If the catalog is password-protected, the update- or higher-level password is required. The catalog must be specified in a JOBCAT or STEPCAT DD statement. If you do not specify the WORKCAT parameter, the catalog will be chosen using existing rules for "Order of Catalog Use: DEFINE" in the chapter "Creating and Cataloging Objects." Further constraints must be followed if you do not specify the WORKCAT parameter, that is, it must be the first DD statement in the JOBCAT or STEPCAT DD statement concatenation. However, if the work file catalog is the master catalog, then you must specify the WORKCAT parameter and it must be the last DD statement in the JOBCAT or STEPCAT DD statement concatenation. The resultant catalog must not be the same one being reset.

Abbreviation: WCAT

WORKFILE(dname[/password])

specifies the name of a DD statement that identifies a VSAM data set name and a list of volume serial numbers of volumes containing VSAM data spaces and units to be used by RESETCAT for a work file. These volumes must be owned by the catalog in which the work file is defined. This data set will be defined by RESETCAT and used as temporary storage while processing the command; it will be deleted at

the end of the command. The data set must not be already defined. Space requirements for the data set are noted in "WORKFILE Space Requirements." If you do not specify WORKFILE, IDCUT1 is used as a default dname.

password

If you desire, you may specify password, a password to be used by RESETCAT to password-protect the work file. Since the work file will contain a copy of the catalog records, it may contain security-sensitive information. The password you specify will become the work file's master password. You may specify the password only if dname is explicitly specified.

Abbreviation: WFILE

RESETCAT EXAMPLES

Resetting a Catalog: Example 1

This example illustrates the use of RESETCAT to reset a catalog (USERCAT1), one of whose owned volumes (VSER01) has been restored. The other volume owned by USERCAT1 is VSER02.

```
//RECOVER JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//STPCAT DD DSN=USERCAT1,DISP=OLD
// DD DSN=AMASTCAT,DISP=OLD
//IDCUT1 DD DSN=WORKFILE,DISP=OLD,
// VOL=SER=231401,UNIT=2314,AMP='AMORG'
//CRAVOL1 DD DISP=OLD,VOL=SER=VSER01,UNIT=2314
//CRAVOL2 DD DISP=OLD,VOL=SER=VSER02,UNIT=2314
//DDCAT DD DSN=USERCAT1,DISP=OLD
//SYSIN DD *
RESETCAT -
  CATALOG(USERCAT1/MASTER DDCAT) -
  CRAFILES((CRAVOL1 ALL) (CRAVOL2 NONE)) -
  WORKCAT(AMASTCAT/MCATMRPW) -
  MASTERPW(MCATMRPW) -
  IGNORE
/*
```

The job control language statements are:

- STEPCAT DD, which makes the user catalog, USERCAT1, and the master catalog, AMASTCAT, available for the step. These catalogs must be available for the catalog to be reset and for the work file to be defined, respectively. Note that when the work file is to be defined in the master catalog, the master catalog must be last in the concatenation and the catalog name is supplied via the WORKCAT parameter.
- IDCUT1 DD, which identifies the data set to be used for the work file.
- CRAVOL1 DD, which identifies and allocates a volume whose catalog recovery area is to be used to reset the catalog.
- CRAVOL2 DD, which identifies and allocates a volume whose catalog recovery area may change as a result of resetting the catalog recovery area on volume VSER01. This volume is not reset in the catalog.
- DDCAT DD, which allocates the catalog to be reset.

The RESETCAT command causes catalog USERCAT1 to be reset from the catalog recovery area on VSER01. The RESETCAT command's parameters are:

- CATALOG, which identifies the catalog to be reset. The master password of the catalog is required to update the catalog. The DD statement for the catalog is also identified.
- CRAFILES, which specifies a list of DD statements which identify the volumes to be used to reset the catalog. ALL specifies that the catalog records in the CRA on the volume specified by the CRAVOL1 DD statement are to be reset (in the catalog). NONE allows the CRAVOL2 DD statement to be specified for a CRA volume which is not to be used for reset, but which may be needed as a result of the reset (for example, a multivolume data set resides on a reset volume and on a non-reset volume).

- WORKCAT, which specifies the name of the catalog in which to define the work file. WORKCAT must be specified if the work file is to be defined in the master catalog.
- IGNORE, which specifies that RESETCAT should force the reset of the catalog despite certain errors which may be encountered during the RESETCAT processing.

Resetting a Catalog: Example 2

This example illustrates the use of RESETCAT to reset a catalog (USERCAT1) after the volume (VSER00) on which the catalog resides has been restored. The volumes owned by USERCAT1, VSER01, and VSER02 have been determined to be out of synchronization with the catalog by running a LISTCRA with the COMPARE option for all volumes owned by USERCAT1.

```
//RECOVER JOB ...
//JOB CAT DD DSN=USERCAT2,DISP=OLD
// DD DSN=USERCAT1,DISP=OLD
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//WORKF DD DSN=WORKFILE,DISP=OLD,UNIT=(2314,2),
// VOL=SER=(231401,231402),AMP='AMORG'
//SYSIN DD *
RESETCAT -
  CATALOG(USERCAT1) -
  CRAVOLS((VSER01 2314)(VSER02 2314)) -
  WORKFILE(WORKF/WRKPSW) -
  NOIGNORE
/*
```

The job control language statements are:

- JOBCAT DD, which makes the user catalogs USERCAT2 and USERCAT1 available for the job. These catalogs must be available for the workfile to be defined and for the catalog to be reset, respectively.
- WORKF DD, which identifies the data set to be used for the work file.

Note: All other required volumes will be dynamically allocated by RESETCAT.

The RESETCAT command causes catalog USERCAT1 to be reset from the catalog recovery areas on VSER01 and VSER02. The RESETCAT parameters are:

- CATALOG, which identifies the catalog to be reset.
- CRAVOLS, which specifies a list of volumes and their device types to be used to reset the catalog.
- WORKFILE, which specifies the DD statement for the work file and a password to be used as the work file's master password.
- NOIGNORE, which specifies that RESETCAT should not force the reset of the catalog if certain errors are encountered.

Resetting a Catalog: Example 3

This three-part example illustrates the use of RESETCAT to reset a catalog (USERCAT1 on volume VSER00) which must be reallocated on a different volume (VSER03). Such a situation may occur when an irreparable physical failure has befallen the catalog. In this example, the catalog recovery area for the damaged catalog volume is assumed to be accessible and valid and thus included for reset in the newly defined catalog. Note that the newly defined catalog bears the same name as the damaged catalog. The other volumes owned by USERCAT1 are VSER01 and VSER02. The catalog is first disconnected, then redefined.

```
//DISC      JOB ...
//STEP1     EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//STEP2     EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//UCATDD    DD  VOL=SER=VSER03,UNIT=3330,DISP=OLD
//CRAVOL1   DD  VOL=SER=VSER01,UNIT=2314,DISP=OLD
//CRAVOL2   DD  VOL=SER=VSER02,UNIT=2314,DISP=OLD
//SYSIN     DD  *
            EXPORT USERCAT1 DISCONNECT

//STEP2     EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//UCATDD    DD  VOL=SER=VSER03,UNIT=3330,DISP=OLD
//CRAVOL1   DD  VOL=SER=VSER01,UNIT=2314,DISP=OLD
//CRAVOL2   DD  VOL=SER=VSER02,UNIT=2314,DISP=OLD
//SYSIN     DD  *
            DEFINE UCAT -
                (NAME (USERCAT1) FILE(UCATDD) VOL(VSER03) -
                 RECOVERABLE CYL(20 1)) -
                 DATA(CYL(10 1)) INDEX(CYL(1))

//STEP3     EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//STEP2CAT DD  DSN=USERCAT1,DISP=OLD
//          DD  DSN=AMASTCAT,DISP=OLD
//IDCUT1    DD  DSN=WORKFILE,DISP=OLD,UNIT=2314,
//          VOL=SER=231401,AMP='AMORG'
//SYSIN     DD  *
            RESETCAT -
                CATALOG(USERCAT1) -
                WORKCAT(AMASTCAT) -
                CRAVOLS((VSER00 3330)(VSER01 2314)(VSER02 2314))
/*
```

STEP1: The user catalog, USERCAT1, is disconnected from the system by using the EXPORT command.

STEP2: The user catalog, USERCAT1, is redefined on the volume VSER03.

STEP3: The user catalog USERCAT1 is reset from the previous catalog's owned volumes.

The job control language statements are:

- STEP2CAT DD, which makes the user catalog, USERCAT1, and the master catalog, AMASTCAT, available for the step. These catalogs must be available for the catalog to be reset, and for the work file to be defined, respectively.
- IDCUT1 DD, which identifies the data set to be used for the work file.

The RESETCAT command causes the redefined catalog USERCAT1 to be reset from the catalog recovery areas on volumes VSER00, VSER01, and VSER02. The RESETCAT command's parameters are:

- CATALOG, which identifies the catalog to be reset.
- WORKCAT, which specifies the name of the catalog in which to define the workfile. WORKCAT must be specified if the work file is to be defined in the master catalog.

- CRAVOLS, which specifies a list of volumes and their device types to be used to reset the catalog.

VERIFY

The format of the VERIFY command is:

VERIFY	{FILE(dname[/password]) DATASET(entryname[/password])}
--------	--

VERIFY can be abbreviated: VFY

VERIFY PARAMETERS

Required Parameters

FILE(dname[/password])

dname specifies the name of a DD statement that identifies the cluster or component to be verified.

password

is the control or master password of a password-protected cluster or component or the master password of a password-protected catalog.

DATASET(entryname[/password])

specifies the name and password of the object to be verified. If DATASET is coded, the object is dynamically allocated.

password

is the control or master password of a password-protected object, or is the master password of the object's (password-protected) catalog.

Abbreviation: DS

Note: When sharing data sets between different processors, it is recommended that you run VERIFY as the first step of a jobstream to prevent job termination caused by an open ACB error code, if the other processor already has the data set open.

When you verify a data set in create mode it is only valid for an ESDS defined with the RECOVERABLE attribute.

VERIFY EXAMPLE

Upgrading a Data Set's End-of-File Information: Example 1

When a data set that was improperly closed (that is, closed as a result of system failure) is opened, the VSAM Open routines set a return code to indicate that the data set's cataloged information might not be accurate. The user can upgrade the EOD and EOKR information (so that it is accurate when the data set is next opened) by closing the data set² and issuing the VERIFY command:

```
//VERIFY JOB ...
//JOB CAT DD DSNAME=D27UCAT1,DISP=SHR
//FIXEOD EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

LISTCAT ENTRIES(FAROUT) -
ALL

VERIFY DATASET(FAROUT)

LISTCAT ENTRIES(FAROUT) -
ALL

/*
```

SYSPRINT is a DD statement required in all Access Method Services job steps to describe the output data set for messages.

The first LISTCAT command lists the data set's cataloged information, showing the data set's parameters as they were when the data set was last closed properly.

The VERIFY command updates the data set's cataloged information to show the data set's real EOD and EOKR values.

The second LISTCAT command lists the data set's cataloged information again. This time, the EOD and EOKR information show the point at which processing stopped due to system failure. This information should help the user determine how much of the data was added correctly before the system failed.

Note: VERIFY will only update the high-used RBA fields for the data set, and not any record counts.

See "Appendix B: Interpreting LISTCAT Output Listings" for details on the order in which catalog records are to be listed and the meanings of the listed records.

² When the data set is opened (first OPEN after system failure), VSAM Open sets a "data set improperly closed" return code. When the data set is closed properly, VSAM Close resets the "data set improperly closed" indicator but does not upgrade erroneous catalog information that resulted from the system failure. Subsequently, when the data set is next opened, its EOD and EOKR information might still be erroneous (until VERIFY is issued to correct it), but VSAM Open sets the "data set opened correctly" return code.

CONTROLLING COMMAND EXECUTION

This chapter describes:

- IF-THEN-ELSE command sequence, which is used to control command execution on the basis of condition codes.
- DO-END command sequence, which specifies more than one functional Access Method Services command and its parameters.
- SET command, which is used to reset condition codes.
- PARM command, which is used to specify diagnostic-aids and printed-output options.

These commands cannot be used when Access Method Services is being executed under TSO.

CONDITION CODES

The condition codes that are tested in the IF-THEN-ELSE command sequence are:

- 0, which indicates that the function was executed as directed and expected. Some informational messages may have been issued.
- 4, which indicates that some problem was met in executing the complete function, but it was possible to continue. The continuation might not provide the user with exactly what is wanted, but no permanent harm will have been done by such continuation. A warning message was issued. An example of the kind of problem encountered is: the system was unable to locate an entry in a LISTCAT command.
- 8, which indicates that a requested function was completed, but major specifics were unavoidably bypassed. For example, an entry to be deleted or altered could not be found in the catalog, or a duplicate name was found while an entry was being defined and the define action was terminated.
- 12, which indicates that the requested function could not be performed. This condition code is set as a result of a logical error. A logical error condition exists when inconsistent parameters are specified, when a too small or too large value is specified for key length, recordsize, or bufferspace, or when required parameters are missing. More information on logical errors that occur during VSAM record processing is in the OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide.
- 16, which indicates that a severe error occurred that caused the remainder of the command stream to be flushed. This condition code might be set as a result, for example, of the following: a system output data set cannot be opened (a SYSPRINT DD statement was missing, for example); an unrecoverable error occurred in a system data set; or Access Method Services encountered improper IF-THEN-ELSE command sequences.

The condition codes that are tested in the IF-THEN-ELSE command sequence and that can be set by the SET command cannot be passed from one job step to the next. However, the final maximum condition code is passed to the OS/VS system when the Access Method Services processor returns control to the system.

IF-THEN-ELSE COMMAND SEQUENCE

The IF-THEN-ELSE command sequence is used to control command execution.

The format of the IF-THEN-ELSE command sequence is:

IF	<pre>{LASTCC MAXCC}[b]comparand[b]number THEN[b]command DO command set END] [ELSE[b]command] DO command set END]]</pre>
-----------	---

where:

IF specifies that one or more functional commands is to be executed based on a test of a condition code. The condition code is set by a SET command or is set to reflect the completion status of previous functional commands.

LASTCC specifies that the condition code value that resulted from the immediately previous function command is to be compared as indicated by the comparand number to determine whether the THEN action is to be performed. See "Condition Codes" earlier in this chapter for the meaning of condition codes.

MAXCC specifies that the maximum condition code value that has been established by any previous function command or by a SET command is to be compared as indicated by the comparand number to determine whether the THEN action is to be performed. See "Condition Codes" earlier in this chapter for the meaning of condition codes.

comparand specifies the comparison to be made between the variable and the following number. This can be any of six possible comparisons:

Equal, specified as '=' or 'EQ'

Not equal, specified as '!=' or 'NE'

Greater than, specified as '>' or 'GT'

Less than, specified as '<' or 'LT'

Greater than or equal, specified as '>=' or 'GE'

Less than or equal, specified as '<=' or 'LE'

number specifies the decimal integer that is to be compared with MAXCC or LASTCC. The number can be up to ten digits long. Values greater than 16 are reduced to 16; both LASTCC and MAXCC are initialized to zero upon entry to Access Method Services. See "Condition Codes" earlier in this chapter for the meaning of condition codes.

THEN specifies that a single command or a group of commands (introduced by DO) are to be executed if the comparison was true. THEN can be followed by another IF command.

ELSE

specifies that a single command or a group of commands (introduced by DO) is to be executed if the previous comparison is false. ELSE can be followed by another IF command.

When an IF command appears in a THEN or ELSE clause, it is called a nested IF command. The maximum level of nesting allowed is 10, starting with the first time you specify IF.

Within a nest of IF commands, the innermost ELSE clause is associated with the innermost THEN clause, the next innermost ELSE clause with the next innermost THEN clause, and so on. (To say it another way, each ELSE is matched with the nearest preceding unmatched THEN.) Should there be an IF command that does not require an ELSE clause, follow the THEN clause with a null ELSE clause (ELSE), unless the nesting structure does not require one.

DO-END Command Sequence

DO

specifies that the group of commands that follow is to be treated as a single unit, that is, to be executed as a result of a single IF command. The set of commands is terminated by END. Commands following a DO must begin on a new line.

END

specifies the end of a set of commands initiated by the nearest unended DO. END must be on a line by itself.

Note: Continuation characters should not be used in the DO-END sequence, as they may be taken as a null command or cause unpredictable results.

NULL COMMANDS

If THEN or ELSE is not followed by a continuation character or by a command in the same record, the THEN or ELSE results in no action.

The null command supports an ELSE command that balances an IF-THEN-ELSE command sequence, and allows null THEN commands. The null command is, in essence, a THEN or ELSE command that is not followed by a command-continuation character.

If you want to indicate a null ELSE command, specify:

```
ELSE
```

If you want to indicate a null THEN command, specify:

```
IF ... THEN  
ELSE ...
```

The null command is used to indicate that no action is to be taken if the IF clause is satisfied (a null THEN command) or if the IF clause is not satisfied (a null ELSE command).

SET COMMAND

The SET command is used to change or reset a previously defined condition code. See "Condition Codes" earlier in this chapter for the meaning of condition codes. It is possible to terminate all processing simply by setting MAXCC or LASTCC to 16.

The format of the SET command is:

SET	(MAXCC LASTCC)[b]=[b]number
-----	-----------------------------

SET

specifies that a condition-code value is to be set. A SET command that follows a THEN or ELSE that is not executed does not cause the value of LASTCC or MAXCC to be altered.

MAXCC

specifies that the value to be reset is the maximum condition code set by a previous functional command. Setting MAXCC does not affect the value of LASTCC.

LASTCC

specifies that the value to be reset is the condition code set by the immediately previous functional command.

number

specifies the value to be assigned to MAXCC or LASTCC. The maximum value that can be assigned is 16; a greater value will be reduced to 16. If the value assigned to LASTCC is greater than the value of MAXCC, MAXCC is set equal to the higher value.

PARM COMMAND

The PARM command specifies processing options to be used during execution. These options remain in effect until changed by another PARM command. You can also specify these options in the PARM field of an EXEC statement (in the JCL).

The format of the PARM command is:

PARM	[TEST([[TRACE] [AREAS(areaid[b areaid...]])] [FULL((dumpid[b count1[b count2]]) [(dumpid...)...]])] OFF)] [GRAPHICS(CHAIN(chain) TABLE(mname))] [MARGINS(leftmargin b rightmargin)]
------	---

where:

TEST(

[[TRACE]
[AREAS(areaid[b areaid...]])]
[FULL((dumpid[b count1[b count2]])
 [(dumpid...)...]])]
OFF)

specifies the diagnostic aids to be used. Once the TEST option has been established, it remains in effect until it is reset by another PARM command. The TRACE, AREAS, and FULL parameters may be used concurrently.

TRACE

specifies that trace tables are to be listed whenever the built-in dump points of the processor are encountered.

AREAS(areaid[b areaid...])

identifies modules that are to have selected variables dumped at their dump points. Each item in the areaid is a two-character area-identifier defined within the implementation. See OS/VS2 Access Method Services Logic for more information.

FULL(({dumpid}[b count1][b count2][dumpid...].]) specifies that a region dump, as well as the trace tables and selected variables, are to be provided at the specified points. dumpid specifies the four-character identifier of the dump point. See OS/VS2 Access Method Services Logic for more information.

count1
is a decimal integer that specifies the number of times (default is 1) the program is to go through the dump point before beginning the dump listing.

count2
specifies a decimal integer which is the number of times (default is 1) through the dump-point that dumps are to be listed.

If the FULL keyword is used, an AMSDUMP DD statement must be provided. For example:

```
//AMSDUMP DD SYSOUT=A
```

OFF
specifies that testing is to stop.

GRAPHICS(CHAIN(chain)|TABLE(mname)) specifies the print chain/train graphic character set or a special graphics table to be used in producing the output. Any character to be printed is translated to the bit pattern found in such a table at the position corresponding to its numeric value (0 through 255). If the print chain does not have a graphic for a byte's bit pattern, the table should specify a period as the output graphic.

CHAIN(AN|HN|PN|QN|RN|SN|TN)
specifies the graphic character set of the print chains or trains you wish employed. PN is used by the processor unless explicitly directed to use another set of graphics.

TABLE(mname)
specifies the name of a user-supplied table. This 256-byte table defines the graphics for each of the 256 possible bit patterns. It must be stored as a module accessible through the LOAD macro (VS).

MARGINS(leftmargin b rightmargin)
specifies that the margins of input records on which command statements are written are to be changed. The normal left and right margins are 2 and 72, respectively. If MARGINS is coded, all subsequent input records are scanned in accord with the new margins. This feature may be used in conjunction with the comment feature: respecification of margins could be used to cause the /* and */ characters to be omitted from the scan and so cause comments to be treated as commands.

leftmargin
specifies the location of the left margin.

rightmargin
specifies the location of the right margin. The right margin value must be greater than the left margin value.

CONTROL COMMAND EXECUTION EXAMPLES

The examples in the topics that follow show the use of the IF-THEN-ELSE command sequence, the SET command, and the PARM command.

Control Command Execution: Example 1

In this example, nested IF commands are used to determine whether a REPRO, DELETE, or PRINT command is to be executed.

```
IF LASTCC > 4 -  
  THEN IF MAXCC < 12 -  
    THEN REPRO...  
    ELSE DELETE...  
  ELSE IF LASTCC = 4 -  
    THEN  
    ELSE PRINT...
```

This example specifies that if the value of LASTCC is greater than 4, then the value of MAXCC is to be tested. If the value of MAXCC is less than 12, the REPRO command is executed, while if the value of MAXCC is 12 or greater, the DELETE command is executed instead. Again, if the value of LASTCC is 4 or less, LASTCC is tested for being exactly 4: no action is to be taken in this case. If, however, LASTCC is less than 4, the PRINT command is to be executed.

Control Command Execution: Example 2

In this example, nested IF commands are used to determine whether a REPRO or PRINT command is executed.

```
IF LASTCC > 4 -  
  THEN IF MAXCC < 12 -  
    THEN REPRO ...  
    ELSE  
  ELSE IF LASTCC = 4 -  
    THEN PRINT ...
```

Should the first IF command determine that LASTCC is greater than 4, and the second IF command determine that MAXCC is 12 or greater, no functional command in the example is executed. The null ELSE command is employed here to specify that the next ELSE is to correspond to the first THEN.

Control Command Execution: Example 3

In this example, if the maximum condition code is zero, an entry from a catalog is listed and a data set is printed.

```
IF MAXCC=0 THEN DO  
  LISTCAT CATALOG (AMASTCAT/MST27) ENT (MN01.0005)  
  PRINT INFILE (AJK006)  
END  
ELSE . . . .
```

Control Command Execution: Example 4

If you want to list a catalog and print a data set if the last condition code is zero, but list its catalog entry before and after a VERIFY command if the last condition code is greater than zero, specify:

```
IF LASTCC = 0 THEN DO
LISTCAT
PRINT INFILE (AJK006)
END
ELSE DO
LISTCAT ENTRY (AJK006) ALL
VERIFY FILE (AJKJCL6)
LISTCAT ENTRY (AJK006) ALL
END
```

Control Command Execution: Example 5

If you want to set the last condition code established to 12, specify:

```
SET LASTCC=12
```

Control Command Execution: Example 6

If you want to replace the highest condition code established in processing so far with 8, specify:

```
SET MAXCC=8
```

APPENDIX A. EXAMPLES OF JOBS USING ACCESS METHOD SERVICES COMMANDS

This set of examples is meant to show the wide range of functions provided by Access Method Services. It is assumed that the VSAM system catalog that exists is recoverable and security protected at the update-, control- and master-password levels.

EXAMPLE 1: DEFINE A VSAM USER CATALOG

In this example, a user catalog is defined. The catalog is password protected. The user catalog in this example is security protected and is not defined with the recoverable attribute. The catalog's data space is defined to be 15 cylinders. The space explicitly specified for the data and index component of the catalog is taken from the catalog's data space. Since the total data space for the data and index component is less than the total catalog data space, the remaining catalog data space is available for suballocation. The volume VSER02 cannot be referenced by other VSAM catalogs since it now belongs to the catalog which it contains.

```
//DEFCAT JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

        DEFINE USERCATALOG -
          (NAME(D27UCAT1) -
           MASTERPW(UCATMRPW) -
           UPDATEPW(UCATUPPW) -
           FOR(365) -
           CYLINDERS(15 5) -
           VOLUMES(VSER02) -
          DATA( -
            CYLINDERS(3 1) ) -
          INDEX( -
            CYLINDERS(1 ) ) -
          CATALOG(AMASTCAT/MCATUPPW)
/*
```

Explanation of Commands

1. The USERCATALOG parameter is required and NAME specifies the user catalog being defined.
2. The MASTERPW parameter specifies the master password for this catalog.
3. The UPDATEPW parameter specifies the update password of this catalog.
4. The FOR parameter specifies the retention period for this file—in this case, one year.
5. The CYLINDERS parameter specifies the amount of space to be allocated to the catalog's data space. A space parameter is required.
6. The VOLUMES parameter is required and specifies the volume containing this catalog. Access Method Services will dynamically allocate the catalog's volume; the volume should be mounted permanently RESIDENT or RESERVED to ensure successful dynamic allocation.
7. The DATA parameter specifies the amount of space to be allocated to the catalog's data component.

8. The INDEX parameter specifies the amount of space to be allocated to the catalog's index component. VSAM adds together the amount of space specified via DATA and INDEX and determines the appropriate proportions for each component.
9. The CATALOG parameter is required since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

EXAMPLE 2: DEFINE A VSAM USER CATALOG AND A VSAM DATA SPACE

This example defines a user catalog by using the previously-defined user catalog, D27UCAT1, as a model. The second part of this example defines a VSAM data space on volume VSER04. Both volumes in this example become owned by the respective catalogs, D27UCAT2 and AMASTCAT; they cannot be referenced by any other VSAM catalogs.

The third part of this example defines alias names for the user catalogs, D27UCAT1 and D27UCAT2.

This example depends on the successful completion of the previous example for the user catalog, D27UCAT1.

```
//DEFCAT2 JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

DEFINE USERCATALOG -
(NAME(D27UCAT2) -
MODEL(D27UCAT1/UCATMRPW D27UCAT1) -
VOLUME(VSER03) -
CYLINDER(15 5)) -
CATALOG(AMASTCAT/MCATUPPW)

DEFINE SPACE -
(VOLUME(VSER04) -
CYLINDERS(10 1) ) -
CATALOG(AMASTCAT/MCATUPPW)

DEFINE ALIAS -
(NAME(D40) -
RELATE(D27UCAT1)) -
CATALOG(AMASTCAT/MCATUPPW)

DEFINE ALIAS -
(NAME(D50) -
RELATE(D27UCAT2)) -
CATALOG(AMASTCAT/MCATUPPW)
/*
```

Explanation of Commands

The first DEFINE command defines a user catalog on volume VSER03. The catalog's attributes are modeled from a previously-defined user catalog, D27UCAT1. The VOLUME and CYLINDERS parameters are required parameters, even through a model catalog is used.

1. The USERCATALOG parameter is required and NAME specifies the name of the user catalog, D27UCAT2.
2. The MODEL parameter names the object to be used as a model and its master password, D27UCAT1/UCATMRPW. The next field names the catalog that contains the model object's catalog entries. In this case, D27UCAT1's self-describing entries describe the object model, D27UCAT1.

3. The VOLUME parameter is required and names the volume that is to contain the user catalog. Access Method Services will dynamically allocate the catalog's volume; the volume should be mounted permanently RESIDENT or RESERVED to ensure successful dynamic allocation.
4. The CYLINDERS parameter specifies the primary and secondary allocation amounts (of cylinders) to be allocated to the catalog's data space. A space parameter specifying the total amount of space is required. The space is allocated to data and index components of the catalog in the same way as that of the model.
5. The CATALOG parameter is required since the master catalog is password protected. It names the master catalog, AMASTCAT, and specifies its update password. When a user-catalog is defined, a user catalog connector entry is built and written into the VSAM master catalog. The user-catalog connector entry enables the user catalog's volume to be located through the master catalog.

The second DEFINE command defines a VSAM data space on volume VSER04 and establishes the volume's ownership by the master catalog, AMASTCAT. All subsequent VSAM objects created on volume VSER04 are required to be cataloged in the volume's catalog, AMASTCAT.

1. The SPACE parameter is required.
2. The VOLUME parameter is required and names the volume that is to contain the data space. Access Method Services will dynamically allocate the volume; the volume should be mounted permanently RESIDENT or RESERVED to ensure successful dynamic allocation.
3. The CYLINDERS parameter specifies the primary and secondary allocation amounts of cylinders to be allocated to the data space. A space parameter is required.
4. The CATALOG parameter identifies the master catalog, AMASTCAT, and supplies its update password.

The defines of the alias names permit references to the respective user catalogs by the alias names. Alias names for VSAM catalogs also provide a way by which data set names can be used to direct the search for catalog information about data sets. Some of the following examples will show how data set naming conventions are used to simplify JCL specification of JOBCAT and STEPCAT DD statements.

1. The ALIAS parameter is required and NAME specifies the alias name of the user catalog.
2. The RELATE parameter specifies the catalog for which the alias is defined.
3. The CATALOG parameter supplies the update password of the master catalog, AMASTCAT.

EXAMPLE 3: DEFINE VSAM DATA SETS

This example defines four VSAM data sets: two key-sequenced data sets into the user catalog, D27UCAT1, a relative-record data set into the system catalog, and a reusable entry-sequenced data set into the user catalog, D27UCAT2.

This example depends on the successful completion of Example 1 for the user catalog.

```

//DEFVSAM JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

```

```

DEFINE -
  CLUSTER ( -
    NAME(D40.MYDATA) -
    VOLUMES(VSER02) -
    RECORDS(1000 500) -
  DATA( -
    NAME(D40.KSDATA) -
    KEYS(15 0) -
    RECORDSIZE(250 250) -
    FREESPACE(20 10) -
    BUFFERSPACE(25000) ) -
  INDEX( -
    NAME(D40.KSINDEX) -
    IMBED ) -
  CATALOG(D27UCAT1/UCATUPPW)

DEFINE CLUSTER -
  (NAME(D40.EXAMPLE.KSDS1) -
   READPW(KSD1PSWD) -
   MODEL(D40.MYDATA)) -
  DATA( -
    (NAME(D40.EXAMPLE.KSDS1.DATA) -
     UNIQUE -
     CYLINDERS(2 1) ) -
  INDEX( -
    (NAME(D40.EXAMPLE.KSDS1.INDEX) -
     UNIQUE -
     CYLINDERS(1 1) ) -
  CATALOG(D27UCAT1/UCATUPPW)

IF LASTCC = 0 -
  THEN -
  LISTCAT ENTRIES( -
    D40.EXAMPLE.KSDS1/KSD1PSWD) -
  ALL

DEFINE CLUSTER -
  (NAME(EXAMPLE.RRDS1) -
   VOLUMES(VSER04) -
   TRACKS(10 5) -
   RECORDSIZE(100 100) -
   NUMBERED ) -
  CATALOG(AMASTCAT/MCATUPPW)

IF LASTCC = 0 -
  THEN -
  LISTCAT ENTRIES( -
    EXAMPLE.RRDS1) -
  CLUSTER -
  ALL

DEFINE CLUSTER -
  (NAME(D50.EXAMPLE.ESDS1) -
   MASTERPW(ESD1MRPW) -
   UPDATEPW(ESD1UPPW) -
   VOLUMES(VSER03) -
   SPANNED -
   REUSE -
   NONINDEXED -
   CYLINDERS(2 1) -
   RECORDSIZE(2500 3000) ) -
  CATALOG(D27UCAT2/UCATMRPW)

IF LASTCC = 0 -
  THEN -
  LISTCAT ENTRIES( -
    D50.EXAMPLE.ESDS1) -
  ALLOCATION

```

```
/*
```

Explanation of Commands

The first DEFINE command defines a key-sequenced data set on volume VSER02. The high-level name of the data set is the alias name of the catalog into which it is being defined.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined.
2. The VOLUMES parameter is required and specifies the volume containing the data set.
3. The RECORDS parameter specifies the space to be allocated to the cluster. This is a required parameter.
4. The DATA parameter is required when attributes are to be explicitly specified for the data component of the cluster. The NAME parameter specifies the name of the data component.
5. The KEYS parameter specifies the key length and offset.
6. The RECORDSIZE parameter specifies the average and maximum record sizes.
7. The BUFFERSPACE parameter is specified for improved performance.
8. The INDEX parameter is required when attributes are to be explicitly specified for the index component of the cluster. The NAME parameter specifies the name of the index component.
9. The IMBED parameter specifies that the index sequence set is to be placed with the data component.
10. The CATALOG parameter is required since the catalog is password protected.

The second DEFINE command defines a unique key-sequenced data set on volume VSER04. The high-level name of the data set is the alias name of the catalog into which it is being defined. This example shows how data set attributes can be specified by modeling and direct specification.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined.
2. The READPW parameter specifies the read password of this cluster. Since no master password is defined, this password will be propagated up to the master level. Thus this cluster is security protected even though its model was not protected.
3. The MODEL parameter specifies the name of the data set to be used as the model.
4. The DATA parameter is required if attributes are to be specified for the data component. The NAME parameter specifies the name of the data component. If a name is not specified, a name is generated.
5. The UNIQUE parameter specifies that this portion of the data set is the only one that occupies the data space allocated to it.
6. The CYLINDERS parameter specifies the amount of space to be allocated to the cluster's data component.
7. The INDEX parameter is required if attributes are to be specified for the index component. The NAME parameter specifies the name of the index component. If a name is not specified, a name is generated.

8. The UNIQUE parameter specifies that this portion of the data set is the only one that occupies the data space allocated to it.
9. The CYLINDERS parameter specifies the amount of space to be allocated to the cluster's index component.
10. The CATALOG parameter is required, since the user catalog is password protected. It specifies the name of the user catalog and its update password which is required to define into a protected catalog.

If the define of the unique key-sequenced data set was successful, then the following LISTCAT command is executed. The high-level name of the data set will direct the LISTCAT to the appropriate user catalog.

1. The ENTRIES and ALL parameters cause the entire catalog description of the data set just defined to be listed.

The third DEFINE command defines a suballocated relative-record data set into the VSAM data space on volume VSER04 which is owned by the master catalog.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined.
2. The VOLUMES parameter is required and specifies the volume containing this data set. Because the master catalog is recoverable, Access Method Services will dynamically allocate this volume to access the catalog recovery area for the data set. This requires that the volume be mounted permanently RESIDENT or RESERVED.
3. The TRACKS parameter specifies the amount of space allocated to this data set. A space parameter is required.
4. The RECORDSIZE parameter specifies the average and maximum record sizes which, in the case of a relative-record data set, must be equal.
5. The NUMBERED parameter is required to specify that this is a relative-record data set.
6. The CATALOG parameter is required since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

If the define of the suballocated relative-record data set was successful, then the following LISTCAT command is executed.

1. The ENTRIES, CLUSTER, and ALL parameters cause the entry just defined to be listed, limited, however, to the cluster entry (that is, the data component's entry is not listed).

The fourth DEFINE command defines a suballocated entry-sequenced data set on volume VSER03.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined. The high-level name of the data set is the alias name of the catalog into which it is being defined.
2. The MASTERPW parameter specifies the master password of this cluster.
3. The UPDATEPW parameter specifies the update password of this cluster.
4. The VOLUMES parameter is required and specifies the volume containing this data set.

5. The SPANNED parameter specifies that records may span control interval boundaries.
6. The REUSE parameter specifies that the data set can be reused, that is, reloaded without being deleted and redefined.
7. The NONINDEXED parameter is required to override the default (INDEXED).
8. The CYLINDERS parameter specifies the amount of space to be allocated to this data set. A space parameter is required.
9. The RECORDSIZE parameter specifies the average and maximum record sizes.
10. The CATALOG parameter specifies the name of the catalog into which the cluster is to be defined. The catalog's update or higher-level password is required.

If the define of the suballocated entry-sequenced data set was successful, then the following LISTCAT command is executed. The high-level name of the data set will direct the LISTCAT to the appropriate catalog.

1. The ENTRIES and ALLOCATION parameters cause the data set entry just defined to be listed, limited, however, to only volume and allocation information.

EXAMPLE 4: DEFINE NONVSAM AND VSAM DATA SETS

This example defines a nonVSAM data set into a user catalog, a VSAM key-sequenced data set into the system catalog, and a VSAM entry-sequenced data set into a user catalog.

```

//DEFVSM2 JOB      ...
//JOB CAT DD      DSN=D27UCAT2,DISP=SHR
//STEP1 EXEC      PGM=IDCAMS
//SYS PRINT DD    SYSOUT=A
//SYS IN DD       *

DEFINE NONVSAM -
    (NAME(EXAMPLE.NONVSAM1) -
     VOLUMES(231401) -
     DEVICETYPES(2314) ) -
    CATALOG(D27UCAT1/UCATUPPW)

IF LASTCC = 0 -
THEN -
LISTCAT NONVSAM -
    ALL -
    CATALOG(D27UCAT1)

DEFINE CLUSTER -
    (NAME(EXAMPLE.KSDS2) ) -
    DATA( -
        NAME(EXAMPLE.KSDS2.DATA) -
        MASTERPW(DAT2MRPW) -
        UPDATEPW(DAT2UPPW) -
        READPW(DAT2RDPW) -
        RECORDS(500 100) -
        EXCEPTIONEXIT(DATEXIT) -
        ERASE -
        FREESPACE(20 10) -
        KEYS(6 4) -
        RECORDSIZE(80 100) -
        VOLUMES(VSER04) ) -
    INDEX( -
        NAME(EXAMPLE.KSDS2.INDEX) -
        MASTERPW(IND2MRPW) -
        UPDATEPW(IND2UPPW) -
        READPW(IND2RDPW) -
        RECORDS(300 300) -
        IMBED -
        VOLUMES(VSER04) ) -
    CATALOG(AMASTCAT/MCATUPPW)

IF LASTCC = 0 -
THEN -
LISTCAT DATA -
    ALL -
    ENTRIES( -
        EXAMPLE.KSDS2/DAT2MRPW) -
    CATALOG(AMASTCAT)

DEFINE CLUSTER -
    (NAME(EXAMPLE.ESDS2) -
     VOLUMES(VSER03) -
     SPANNED -
     CYLINDERS(2 1) -
     NONINDEXED -
     REUSE -
     MASTERPW(ESD2MRPW) -
     CONTROLPW(ESD2CTPW) -
     UPDATEPW(ESD2UPPW) -
     READPW(ESD2RDPW)) -
    CATALOG(D27UCAT2/UCATMRPW)

IF LASTCC = 0 -
THEN -
DO -
LISTCAT ENTRIES( -
    EXAMPLE.ESDS2/ESD2MRPW) -
    ALL
LISTCAT NAME -
    CATALOG(AMASTCAT/MCATMRPW)
END

```

/*

Explanation of Job Control Language Statements

1. The JOBCAT DD statement describes the user catalog D27UCAT2 as a job catalog. All references will be to this job catalog unless otherwise directed.

Explanation of Commands

The first DEFINE command defines an existing nonVSAM data set into user catalog D27UCAT1.

1. The NONVSAM parameter is required and NAME specifies the nonVSAM object being defined.
2. The VOLUMES parameter is required and specifies the volume containing the data set.
3. The DEVICETYPES parameter is required and specifies the device type of the volume.
4. The CATALOG parameter specifying the name of the user catalog is required because:
 - A job catalog also appears in the job control language, so this parameter explicitly directs the define to the user catalog.
 - The catalog is protected and its update password is required for the define.

If the definition of the nonVSAM entry was successful, then the following LISTCAT command is executed.

1. The NONVSAM and ALL parameters cause all the nonVSAM entries cataloged in D27UCAT1 to be listed.
2. The CATALOG parameter directs the LISTCAT to a specific user catalog.

The second DEFINE command defines a key-sequenced data set into a VSAM data space owned by the master catalog. Note that attributes are specified at the data and index level rather than at the cluster level. Access Method Services will dynamically allocate the volume containing the catalog recovery area for the cluster, since the system catalog is recoverable. This requires that the volume containing the index component of the cluster be mounted permanently RESIDENT or RESERVED to ensure successful dynamic allocation.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined.
2. The DATA component is explicitly named via the NAME parameter.
3. The MASTERPW, UPDATEPW, and READPW parameters specify the master, update, and read passwords, respectively, of this data component.
4. The RECORDS parameter specifies the amount of space to be allocated to the data component. A space allocation parameter is required.
5. The EXCEPTIONEXIT parameter specifies the name of the routine to be given control if an exception occurs while processing the data component.
6. The ERASE parameter specifies that the data component is to be overwritten with binary zeros when it is deleted.
7. The FREESPACE parameter specifies the percentage of space within control intervals and control areas, respectively, that is to remain free.

8. The KEYS parameter specifies the key length and offset.
9. The RECORDSIZE parameter specifies the average and maximum record sizes.
10. The VOLUMES parameter is required and specifies the volume containing this data component.
11. The INDEX component is explicitly named via the NAME parameter.
12. The MASTERPW, UPDATEPW, and READPW parameters specify the master, update, and read passwords, respectively, of this index component.
13. The RECORDS parameter specifies the amount of space to be allocated to the index component. A space allocation parameter is required.
14. The IMBED parameter specifies that the index sequence set is to be placed with the data component.
15. The VOLUMES parameter is required and specifies the volume containing this index component.
16. The CATALOG parameter is required, since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog. This parameter is also required because a DD statement for the job catalog (JOB CAT) appears in the job control language and this define must, therefore, be explicitly directed to the master catalog.

If the definition of the key-sequenced data set was successful, then the following LISTCAT command is executed.

1. The DATA, ALL, and ENTRIES parameters cause all of the information contained in the data component entry to be listed.
2. The CATALOG parameter directs the LISTCAT to the master catalog.

The third DEFINE command defines an entry-sequenced data set.

1. The CLUSTER parameter is required and NAME specifies the cluster being defined.
2. The VOLUMES parameter specifies the volume (VSER03) which is to contain the data set being defined.
3. The SPANNED parameter specifies that records may span control interval boundaries.
4. The CYLINDERS parameter specifies the amount of space to be allocated to this data set. A space parameter is required.
5. The NONINDEXED parameter is required to override the default (INDEXED).
6. The REUSE parameter specifies that the data set can be reused, that is, reloaded without being deleted and redefined.
7. The MASTERPW, CONTROLPW, UPDATEPW, and READPW parameters specify passwords different from the passwords specified for the data set being modeled.
8. The CATALOG parameter is required, since the user catalog is password protected.

If the define of the entry-sequenced data set was successful, then the following LISTCAT commands are executed.

1. The ENTRIES and ALL parameters of the first LISTCAT command cause all the cataloged information in the entry just defined to be listed.
2. The NAME parameter of the second LISTCAT command causes only the names of the objects cataloged in the master catalog to be listed.
3. The CATALOG parameter specifies the master password of the master catalog which allows access to all objects in the catalog and directs the LISTCAT to the master catalog.

EXAMPLE 5: COPYING AND PRINTING

This example shows various techniques which can be used to load and print data sets using the REPRO and PRINT commands. This example depends on the successful completion of Examples 3 and 4 for the existence of VSAM data sets to load records into. This example also requires that various nonVSAM data sets exist, so that data records can be loaded into the VSAM data sets.

```

//COPYPRNT JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//* NONVSAM INDEXED-SEQUENTIAL DATA SET
//INDSET1 DD DSN=D40.IF100,DISP=OLD,DCB=DSORG=IS,
// VOL=SER=231401,UNIT=2314
//* NONVSAM SEQUENTIAL DATA SET (VARIABLE-LENGTH
//* RECORDS)
//INDSET2 DD DSN=D40.SVB200,DISP=OLD,
// VOL=SER=231401,UNIT=2314
//* NONVSAM SEQUENTIAL DATA SET (FIXED-LENGTH
//* RECORDS)
//INDSET3 DD DSN=D40.SF100,DISP=OLD,
// VOL=SER=231401,UNIT=2314
//* NONVSAM SEQUENTIAL DATA SET
//INDSET4 DD DSN=EXAMPLE.NONVSAM1,DISP=OLD,
// VOL=SER=231401,UNIT=2314
//SYSIN DD *

```

```

/* LOAD A VSAM KEY-SEQUENCED DATA SET */
/* FROM AN ISAM DATA SET */

```

```

REPRO INFILE(INDSET1) -
      OUTDATASET(D40.EXAMPLE.KSDS1/KSD1PSWD)

```

```

IF MAXCC = 0 -
  THEN -
    PRINT INDATASET(D40.EXAMPLE.KSDS1/KSD1PSWD) -
          FROMKEY(X'40F0F0F0F0F6')

```

```

/* LOAD A VSAM ENTRY-SEQUENCED DATA SET FROM AN */
/* EXISTING VARIABLE UNBLOCKED SAM DATA SET */

```

```

REPRO INFILE(INDSET2) -
      OUTDATASET(D50.EXAMPLE.ESDS1/ESD1UPPW) -
      COUNT(030)

```

```

IF LASTCC = 0 -
  THEN -
    PRINT INDATASET(D50.EXAMPLE.ESDS1) -
          FROMADDRESS(170) -
          HEX

```

```

/* LOAD A VSAM RELATIVE-RECORD DATA SET FROM AN */
/* EXISTING SAM DATA SET */

```

```

REPRO INFILE(INDSET3) -
      OUTDATASET(EXAMPLE.RRDS1) -
      SKIP(10)

```

```

IF LASTCC = 0 -
  THEN -
    PRINT INDATASET(EXAMPLE.RRDS1) -
          TONUMBER(25)

```

```

/* PRINT THE CONTENTS OF THE SAM DATA SET */

```

```

PRINT INFILE(INDSET3) -
      COUNT(20) -
      CHARACTER

```

```

/* LOAD A VSAM KEY-SEQUENCED DATA SET */
/* FROM A NONVSAM DATA SET */

```

```

REPRO INFILE(INDSET4) -
      OUTDATASET(EXAMPLE.KSDS2)

```

```

IF LASTCC = 0 -
  THEN -
    PRINT INDATASET(EXAMPLE.KSDS2) -
          FROMKEY(AAAAJA) -
          TOKEY(AAAAJ9)

```

```

/*

```

Explanation of Job Control Language Statements

1. The INDSET1 DD statement describes the ISAM data set which is to be copied into the VSAM data set D40.EXAMPLE.KSDS1.
2. The INDSET2 DD statement describes the variable-length SAM data set which is to be copied into the VSAM data set D50.EXAMPLE.ESDS1.
3. The INDSET3 DD statement describes the SAM data set which is to be copied into the VSAM data set EXAMPLE.RRDS1.
4. The INDSET4 DD statement describes the SAM data set which is to be copied into the VSAM data set EXAMPLE.KSDS2.

Explanation of Commands

The first REPRO command causes a VSAM key-sequenced data set to be loaded from an ISAM data set.

1. The INFILE parameter identifies the data set containing the source data. (**Note:** Either the INFILE or INDATASET is required). The dname of the DD statement for this data set must be identical to this name.
2. The OUTDATASET parameter identifies the name of the data set to be loaded. (**Note:** Either the OUTFILE or OUTDATASET is required). The data set is dynamically allocated by Access Method Services. The catalog containing the data set is located by its alias 'D40'. The update or higher password of the VSAM data set is required.

If the REPRO operation was successfully executed, then the contents of the VSAM key-sequenced data set just loaded are printed. The format of the listing is DUMP, since the default is taken.

1. The INDATASET parameter identifies the name of the data set to be printed. (**Note:** Either INFILE or INDATASET is required.) The data set is dynamically allocated by Access Method Services. The catalog containing the data set is located by its alias 'D40'. The update- or higher-level password of the VSAM data set is required.
2. The FROMKEY parameter specifies that printing is to begin with the record whose key (high-order three bytes) is greater than or equal to 'b00006' (that is, the character equivalent of X'40F0F0F0F0F6').

The second REPRO command causes a VSAM entry-sequenced data set to be loaded from an existing variable unblocked SAM data set.

1. The INFILE parameter identifies the data set containing the source data. The dname of the DD statement for this data set must be identical to this name.
2. The OUTDATASET parameter identifies the name of the data set to be loaded. The data set is dynamically allocated by Access Method Services. The catalog containing the data set is located by its alias 'D40'. The update or higher password of the VSAM data set is required.
3. The COUNT parameter specifies that the first 30 records of the SAM data set are to be loaded.

If the REPRO operation was successfully executed, then the contents of the VSAM entry-sequenced data set just loaded are printed in hexadecimal format.

1. The INDATASET parameter identifies the name of the data set to be printed. The data set is dynamically allocated by Access Method Services. The catalog containing the data set is located by its alias 'D40'. Since the data set is not read protected, no password is required.
2. The FROMADDRESS parameter specifies that the first record printed is that record whose relative byte address is exactly equal to 170.
3. The HEX parameter specifies that the listing is to be in hexadecimal format.

The third REPRO command causes a VSAM relative-record data set to be loaded from an existing fixed SAM data set. The relative-record data set can receive only fixed length records that equal its defined record length.

1. The INFILE parameter identifies the data set containing the source data. The dname of the DD statement for this data set must be identical to this name.
2. The OUTDATASET parameter identifies the name of the data set to be loaded. The data set is dynamically allocated by Access Method Services. The data set is cataloged in the master catalog; therefore, no JOBCAT or STEPCAT DD statement is required.
3. The SKIP parameter specifies that the first 10 records of the SAM data set are to be bypassed.

If the REPRO operation was successfully executed, then the contents of the VSAM relative-record data set just loaded are printed.

1. The INDATASET parameter identifies the name of the data set to be printed. The data set is dynamically allocated by Access Method Services. The data set is cataloged in the master catalog; therefore, no JOBCAT or STEPCAT DD statement is required. Since the data set is not read protected, no password is required.
2. The TONUMBER parameter limits the output to those relative records with relative-record number less than or equal to 25.

The PRINT command causes the contents of the SAM data set to be printed.

1. The INFILE parameter identifies the data set to be printed. The dname of the DD statement for this data set must be identical to this name.
2. The COUNT parameter specifies that only 20 records are to be printed.
3. The CHARACTER parameter specifies that the listing is to be in character format.

The last REPRO command causes a VSAM key-sequenced data set to be loaded from a nonVSAM data set.

1. The INFILE parameter identifies the data set containing the source data. The dname of the DD statement for this data set must be identical to this name.
2. The OUTDATASET parameter identifies the name of the data set to be loaded. The data set is dynamically allocated by Access Method Services. The data set is cataloged in the master catalog. Note that a password is not required since the cluster component is not password protected although its data and index components are.

If the REPRO operation was successfully executed, then the contents of the VSAM key-sequenced data set just loaded are printed. The FROMKEY and TOKEY parameters are used to limit the output to a specific range of keys.

1. The INDATASET parameter identifies the name of the data set to be printed. The data set is dynamically allocated by Access Method Services. The data set is cataloged in the master catalog. No password is required, because the cluster component is not password protected.
2. The FROMKEY and TOKEY parameters specify the keys at which printing is to begin and end, respectively.

EXAMPLE 6: RECORD REPLACEMENT

This example shows techniques for modifying the contents of VSAM data sets using the REPRO command.

This example depends on the successful completion of Example 5 for the existence of nonempty VSAM data sets.

```
//REPLACE JOB      ...
//STEP1  EXEC     PGM=IDCAMS
//SYSPRINT DD      SYSOUT=A
//INDSET4 DD      DSNAME=EXAMPLE.NONVSAM2,DISP=OLD,
//          VOL=SER=231401,UNIT=2314
//SYSIN   DD      *

      REPRO      INFILE(INDSET4) -
                OUTDATASET(EXAMPLE.KSDS2) -
                REPLACE

      IF LASTCC = 0 -
        THEN -
          PRINT   INDATASET(EXAMPLE.KSDS2) -
                FROMKEY(AAAAJA) -
                TOKEY(AAAAJ9)

      REPRO      INDATASET(EXAMPLE.RRDS1) -
                OUTDATASET(D50.EXAMPLE.ESDS1/ESD1UPPW) -
                REUSE

      IF LASTCC = 0 -
        THEN -
          PRINT   INDATASET(D50.EXAMPLE.ESDS1)
/*
```

Explanation of Job Control Language Statements

1. The INDSET4 DD statement describes the nonVSAM data set to be copied into the VSAM data set EXAMPLE.KSDS2.

Explanation of Commands

The first REPRO command causes records in the VSAM key-sequenced data set to be replaced with input from a nonVSAM data set.

1. The INFILE parameter identifies the data set containing the source data. (Note: Either INFILE or INDATASET is required.) The dname of the DD statement for this data set must be identical to this name.
2. The OUTDATASET parameter identifies the target data set. (Note: Either OUTFILE or OUTDATASET is required.) Access Method Services will dynamically allocate the data set. The data set is cataloged in the master catalog; therefore, no JOBCAT or STEPCAT DD statement is required. Since the cluster is not password protected, no password is required.

3. The REPLACE parameter causes a record in the output data set having the same key as a record in the input data set to be replaced. Records in the input data set whose key is not already contained in the output data set will be inserted in the output data set.

If the REPRO operation was successfully executed, then the contents of the VSAM key-sequenced data set just changed are printed.

1. The INDATASET parameter identifies the data set to be printed. (Note: Either INFILE or INDATASET is required.) Access Method Services will dynamically allocate the data set. The data set is cataloged in the master catalog; therefore, no JOBCAT or STEPCAT DD statement is required. Since the cluster is not password protected, no password is required.
2. The FROMKEY and TOKEY parameters specify the keys at which printing is to begin and end, respectively.

The second REPRO command causes the VSAM entry-sequenced data set to be loaded from the VSAM relative-record data set.

1. The INDATASET parameter identifies the source data set. Access Method Services will dynamically allocate the data set. The data set is cataloged in the master catalog. Since the cluster is unprotected, no password is required.
2. The OUTDATASET parameter identifies the target data set. Access Method Services will dynamically allocate the data set. The catalog containing the data set is located by its alias 'D50'. The update- or higher-level password is required to load the data set.
3. The REUSE parameter specifies that any records already in the entry-sequenced data set output are to be overwritten since the entry-sequenced data set was defined with the REUSE attribute.

If the REPRO operation was successfully executed, then the entire contents of the reloaded VSAM entry-sequenced data set are printed.

- The INDATASET parameter identifies the data set to be printed. Access Method Services will dynamically allocate the data set. The catalog containing the data set is located by its alias 'D50'. Since no read password exists for this data set, no password is required.

EXAMPLE 7: ALTER THE CATALOGED ATTRIBUTES OF VSAM DATA SETS

This example shows how the cataloged attributes of two VSAM data sets are modified. Each ALTER command is followed by a LISTCAT command, which will execute only if its previous ALTER command completed successfully. The LISTCAT command prints the updated catalog entry.

This example depends on the successful completion of Example 3, which defined the two VSAM data sets whose attributes are being altered.

```

//ALTER      JOB      ...
//STEP1     EXEC PGM=IDCAMS
//SYSPRINT  DD      SYSOUT=A
//SYSIN     DD      *

ALTER -
      D40.EXAMPLE.KSDS1.DATA -
      FREESPACE(10 10)

IF LASTCC = 0 -
  THEN -
    LISTCAT -
      ENTRIES(D40.EXAMPLE.KSDS1.DATA/KSD1PSWD) -
      ALL

ALTER -
      D50.EXAMPLE.ESDS1/ESD1MRPW -
      MASTERPW(ESD1PWMR) -
      CONTROLPW(ESD1PWCT) -
      UPDATEPW(ESD1PWUP) -
      READPW(ESD1PWRD)

IF LASTCC = 0 -
  THEN -
    LISTCAT -
      ENTRIES(D50.EXAMPLE.ESDS1/ESD1PWMR) -
      CLUSTER -
      ALL
/*

```

Explanation of Commands

The first ALTER command shows how a data set's space management attributes are "tuned" for optimum performance.

The data component, D40.EXAMPLE.KSDS1.DATA, of a key-sequenced VSAM data set, D40.EXAMPLE.KSDS1, was defined with 40% free space in both control intervals and control areas. Now that data records have been loaded into the data set, its free space attributes no longer appear to require 40% freespace. It is now desirable to have 10% free space in both control intervals and control areas. The catalog containing the component to be altered is located through its alias 'D40', which is the high-level qualifier of the component name.

- D40.EXAMPLE.KSDS1.DATA names the entry whose attributes are to be altered with this command. No password is required since the data object is not password protected.
- The FREESPACE parameter respecifies percentages of free space that apply to the data component.

The ALTER command is followed by a modal command that examines the condition code set when the ALTER command completes. If the ALTER command completes successfully, the LISTCAT command that immediately follows it prints the changed catalog entry.

- The ENTRIES and ALL parameters explicitly name the entry to be listed and specify that the entire entry is to be listed. Since all information about the data component is to be listed, some information in the associated cluster must be accessed; this requires the cluster's read password.

The second ALTER command shows how a data-set's passwords can be modified. You could use the same technique to provide passwords for an existing VSAM object that does not have passwords.

- D50.EXAMPLE.ESDS1 names the entry whose attributes are to be altered with this command. The entry's master password, ESD1MRPW, is supplied to allow the command to alter the entry's passwords. The catalog containing the entry to be altered is located through its alias 'D50'.

- The MASTERPW, CONTROLPW, UPDATEPW, and READPW parameters respecify passwords for the data set.

If the ALTER command was successful, the LISTCAT command lists all of the cluster entry.

EXAMPLE 8: CREATING AN ALTERNATE INDEX AND ITS PATH

This example defines an alternate index over a previously loaded VSAM key-sequenced base cluster, defines a path over the alternate index to provide a means for processing the base through the alternate index, and builds the alternate index. The alternate index, path, and base cluster must all be defined in the same catalog, in this case, the master catalog. Since the master catalog is recoverable, Access Method Services will dynamically allocate the volume containing the base cluster's index component in order to access its recovery area.

This example depends on the successful completion of Examples 4, 5, and 6 for the existence of the nonempty VSAM base cluster EXAMPLE.KSDS2.

```
//MAKEAIX JOB    ...
//STEP1  EXEC   PGM=IDCAMS
//SYSPRINT DD   SYSOUT=A
//IDCUT1  DD    DSNAME=SORT.WORK.ONE,DISP=OLD,
//          AMP='AMORG',VOL=SER=VSER01,UNIT=2314
//IDCUT2  DD    DSNAME=SORT.WORK.TWO,DISP=OLD,
//          AMP='AMORG',VOL=SER=VSER01,UNIT=2314
//SYSIN   DD    *

        DEFINE ALTERNATEINDEX -
            (NAME(EXAMPLE.AIX) -
             RELATE(EXAMPLE.KSDS2) -
             MASTERPW(AIXMRPW) -
             UPDATEPW(AIXUPPW) -
             KEYS(3 0) -
             RECORDSIZE(40 50) -
             VOLUMES(VSER04) -
             CYLINDERS(2 1) -
             NONUNIQUEKEY -
             UPGRADE ) -
            CATALOG(AMASTCAT/MCATUPPW)

        DEFINE PATH -
            (NAME(EXAMPLE.PATH) -
             PATHENTRY(EXAMPLE.AIX/AIXMRPW) -
             READPW(PATHRDPW)) -
            CATALOG(AMASTCAT/MCATUPPW)

        BLDINDEX INDATASET(EXAMPLE.KSDS2) -
            OUTDATASET(EXAMPLE.AIX/AIXUPPW) -
            CATALOG(AMASTCAT/MCATMRPW)

        PRINT INDATASET(EXAMPLE.PATH/PATHRDPW)
/*
```

Explanation of Job Control Language Statements

- The IDCUT1 and IDCUT2 DD statements describe the dsnames and a volume containing VSAM data space to be made available to BLDINDEX for defining and using two sort work data sets in the event an external sort is performed. This data space will not be used by BLDINDEX if enough virtual storage is available to perform an internal sort.

Explanation of Commands

The first DEFINE command creates a VSAM alternate index over the base cluster EXAMPLE.KSDS2.

1. The NAME parameter is required and names the object being defined.
2. The RELATE parameter is required and specifies the name of the base cluster over which the alternate index is defined.
3. The MASTERPW and UPDATEPW parameters specify the master and update passwords, respectively, for the alternate index.
4. The KEYS parameter specifies the length of the alternate key and its offset in the base cluster record.
5. The RECORDSIZE parameter specifies the length of the alternate-index record. It must be large enough to contain the prime keys for all occurrences of any one alternate key since the alternate index is being defined with the NONUNIQUEKEY attribute.
6. The VOLUMES parameter is required and specifies the volume containing the alternate index EXAMPLE.AIX.
7. The CYLINDERS parameter specifies the amount of space to be allocated to the alternate index. A space parameter is required.
8. The NONUNIQUEKEY parameter specifies that the base cluster can contain multiple occurrences of any one alternate key.
9. The UPGRADE parameter specifies that the alternate index is to reflect all changes made to the base cluster records, for example, additions or deletions of base cluster records.
10. The CATALOG parameter is required, since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

The second DEFINE command defines a path over the alternate index. After the alternate index has been built, opening with the path name will cause processing of the base cluster via the alternate index.

1. The NAME parameter is required and names the object being defined.
2. The PATHENTRY parameter is required and specifies the name of the alternate index over which the path is defined and its master password.
3. The READPW parameter specifies a read password for the path; it will be propagated to master-password level.
4. The CATALOG parameter is required, since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

The BLDINDEX command builds an alternate index. The assumption is made that enough virtual storage will be available to perform an internal sort. However, note that DD statements with the default dnames of IDCUT1 and IDCUT2 have been provided for two external sort work data sets in the event that the assumption is incorrect and an external sort must be performed.

1. The INDATASET parameter identifies the base cluster. Access Method Services will dynamically allocate the base cluster. The base-cluster's cluster entry is not password protected even though its data and index components are.

2. The OUTDATASET parameter identifies the alternate index. Access Method Services will dynamically allocate the alternate index. The update-or higher-level password of the alternate index is required.
3. The CATALOG parameter specifies the name of the master catalog. If it is necessary for BLDINDEX to use external sort work data sets, they will be defined in and deleted from the master catalog. The master password will permit these actions.

The PRINT command causes the base cluster to be printed by means of the alternate key using the path defined to create this relationship.

- The INDATASET parameter identifies the path object. Access Method Services will dynamically allocate the path. The read password of the path is required.

EXAMPLE 9: EXPORTING VSAM DATA SETS

This example shows various methods of exporting data sets to provide backup and portability.

The example depends on the successful completion of previous examples for the existence of the various objects to be exported.

```

//EXPORT JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//RECEIVE DD DSNAME=PORTABLE.TAPE1,LABEL=(1,SL),
// UNIT=2400-3,DISP=(NEW,PASS),VOL=SER=TAPE01,
// DCB=(DEN=3,BLKSIZE=6000)
//SYSIN DD *

EXPORT -
D40.EXAMPLE.KSDS1/KSD1PSWD -
PURGE -
OUTFILE(RECEIVE)

//STEP2 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//RECEIVE DD DSNAME=PORTABLE.TAPE2,LABEL=(2,SL),
// UNIT=2400-3,DISP=(NEW,PASS),VOL=SER=TAPE01,
// DCB=DEN=3
//SYSIN DD *

EXPORT -
D50.EXAMPLE.ESDS1/ESD1PWMR -
OUTFILE(RECEIVE) -
TEMPORARY -
INHIBITSOURCE -
INHIBITTARGET

//STEP3 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//RECEIVE DD DSNAME=PORTABLE.TAPE3,LABEL=(3,SL),
// DISP=(NEW,PASS),VOL=SER=TAPE01,UNIT=2400-3,
// DCB=DEN=3
//SYSIN DD *

EXPORT -
EXAMPLE.AIX/MCATMRPW -
OUTFILE(RECEIVE)

//STEP4 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//RECEIVE DD DSNAME=PORTABLE.TAPE4,LABEL=(4,SL),
// DISP=(NEW,PASS),VOL=SER=TAPE01,UNIT=2400-3,
// DCB=DEN=3
//SYSIN DD *

EXPORT -
EXAMPLE.KSDS2/MCATMRPW -
OUTFILE(RECEIVE)

/*

```

Explanation of Job Control Language Statements

1. The RECEIVE DD statements describe the portable data sets. The record format (VBS) and logical record length are set by EXPORT. The blocksize is set by EXPORT as 2048 except where overridden as in STEP1.

Explanation of Commands

The first EXPORT command causes a key-sequenced VSAM data set to be exported from a user catalog. When it has been exported to a magnetic tape file, the key-sequenced data set is deleted from the user catalog.

Note: When an object is exported, the record format of its records on the portable file is "VBS," and the EXPORT process determines the appropriate record size. However, the RECEIVE DD statement specifies a blocksize (BLKSIZE=6000) to override the blocksize used by the EXPORT process (2048 bytes) and to improve performance.

1. D40.EXAMPLE.KSDS1 names the key-sequenced VSAM data set being exported. Its password, KSD1PSWD, is also supplied (when the data set was defined, its read password propagated upward and all passwords for the data set are KSD1PSWD). Access Method Services will dynamically allocate the cluster. The catalog containing the cluster is located through its alias 'D40', which is the high-level qualifier of the cluster name.
2. The PURGE parameter is required, because the data set was defined with a retention period of 365 days. The data set cannot be exported permanently (that is, deleted from the catalog after its copy is made in the portable file) unless the PURGE parameter is specified to override its cataloged retention period.
3. The OUTFILE parameter names the DD statement that describes and allocates the first file on a magnetic tape reel. (Note: Either OUTFILE or OUTDATASET is required.)

The second EXPORT command causes an entry-sequenced VSAM data set to be exported from a user catalog. When it has been exported to a magnetic tape file, the data set's entry in the user catalog is marked "temporary export" and "inhibit update," which prevents the data set from being modified in any way. The user's program can only read the data set's records. In addition, when the data set's copy on magnetic tape is imported into another system catalog, the data set's entry in the new, or "target," catalog is marked "inhibit update."

1. D50.EXAMPLE.ESDS1 names the entry-sequenced data set being exported. Its master password, ESD1PWMR, is also supplied. (Example 7 shows how the data set's passwords were changed.) Access Method Services will dynamically allocate the cluster. The catalog containing the cluster is located through its alias 'D50', which is the high-level qualifier of the cluster name.
2. The OUTFILE parameter names the DD statement that describes and allocates the second file on a magnetic tape reel.
3. The TEMPORARY parameter specifies that the data set is not to be deleted from the catalog when it is exported.
4. The INHIBITSOURCE parameter specifies that the data set that remains in the "source" catalog and system is not to be updated or modified.
5. The INHIBITTARGET parameter specifies that the data set's exported copy is not to be updated or modified when it has been imported into the "target" catalog and system.

The third and fourth EXPORT commands cause the alternate index and base cluster to be exported from the master catalog. Any paths defined over either object will be exported with their PATHENTRY object. Since the export is permanent, both the base cluster and alternate index will be deleted from the catalog. The alternate index must be exported first, because the delete of the base cluster will cause deletion of all objects defined over it.

The third EXPORT command causes an alternate index to be exported from the master catalog.

1. The name of the alternate index being exported is required. A master password is required for the deletion and to allow VSAM locates against both the alternate index and path to obtain the catalog information (including passwords) to be exported. The master password of the catalog covers all requirements. Access Method Services will dynamically allocate the alternate index; this will also provide access to the catalog recovery area.

2. The `OUTFILE` parameter names the DD statement that describes and allocates the third file on a magnetic tape reel.

The fourth `EXPORT` command causes a base cluster to be exported from the master catalog.

1. The name of the base cluster being exported is required. Since the cluster level is not protected, no password would be required for the deletion. However, a password is required for the VSAM locates against the data and index components to obtain the catalog information (including passwords) to be exported. Since only one password can be supplied, it must be that of the master catalog. Access Method Services will dynamically allocate the cluster; this will also provide access to the catalog recovery area.
2. The `OUTFILE` parameter names the DD statement that describes and allocates the fourth file on a magnetic tape reel.

EXAMPLE 10: IMPORTING VSAM DATA SETS

This example shows various methods of importing data sets.

This example depends on the successful completion of Example 9, which created a portable magnetic tape that contains a copy of each VSAM data set to be imported.

```

//IMPORT  JOB  ...
//STEP1  EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SOURCE  DD  DSNNAME=PORTABLE.TAPE4,LABEL=(4,SL),
//  DISP=(OLD,PASS),VOL=SER=TAPE01,UNIT=2400-3,DCB=DEN=3
//SYSIN   DD  *

      IMPORT -
      INFILE(SOURCE) -
      OUTDATASET(EXAMPLE.KSDS2) -
      CATALOG(AMASTCAT/MCATUPPW) -

//STEP2  EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SOURCE  DD  DSNNAME=PORTABLE.TAPE3,LABEL=(3,SL),
//  DISP=(OLD,PASS),VOL=SER=TAPE01,UNIT=2400-3,DCB=DEN=3
//SYSIN   DD  *

      IMPORT -
      INFILE(SOURCE) -
      OUTDATASET(EXAMPLE.AIX) -
      CATALOG(AMASTCAT/MCATUPPW)

//STEP3  EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SOURCE  DD  DSNNAME=PORTABLE.TAPE2,LABEL=(2,SL),
//  UNIT=2400-3,DISP=(OLD,PASS),VOL=SER=TAPE01,
//  DCB=(DEN=3,LRECL=3004)
//RECEIVE DD  DSN=D50.EXAMPLE.ESDS1,DISP=OLD
//SYSIN   DD  *

      IMPORT -
      INFILE(SOURCE) -
      OUTFILE(RECEIVE) -
      CATALOG(D27UCAT2/UCATMRPW)

//STEP4  EXEC  PGM=IDCAMS
//SYSPRINT DD  SYSOUT=A
//SOURCE  DD  DSNNAME=PORTABLE.TAPE1,LABEL=(1,SL),
//  UNIT=2400-3,DISP=(OLD,PASS),VOL=SER=TAPE01,
//  DCB=(DEN=3,LRECL=479,BLKSIZE=6000)
//SYSIN   DD  *

      IMPORT -
      INFILE(SOURCE) -
      OUTDATASET(D50.EXAMNEW.KSDS1) -
      CATALOG(D27UCAT2/UCATUPPW) -
      OBJECTS( -
        (D40.EXAMPLE.KSDS1 -
          VOLUMES(VSER03) -
          NEWNAME(D50.EXAMNEW.KSDS1) ) -
        (D40.EXAMPLE.KSDS1.DATA -
          NEWNAME(D50.EXAMNEW.KSDS1.DATA) ) -
        (D40.EXAMPLE.KSDS1.INDEX -
          NEWNAME(D50.EXAMNEW.KSDS1.INDEX) ) )
/*

```

Explanation of Job Control Language Statements

- The SOURCE DD statements describe the portable data sets. For standard labeled tapes, the DCB parameters of LRECL and BLKSIZE need not be specified. However, the example shows parameters which would be required for a nonlabeled tape. Unless overridden by DCB parameters, the EXPORT command sets the blocksize to 2048 and the record length to blocksize minus 4.
- The RECEIVE DD statement in STEP3 describes and allocates D50.EXAMPLE.ESDS1, which was exported with the TEMPORARY attribute.

Explanation of Commands

The first and second IMPORT commands will import the base cluster and alternate index exported in Example 9 into the master catalog. The importation causes each component to be newly defined. Since the alternate index cannot be defined until the base cluster has been defined, the base cluster must be imported first. The importation will cause any paths over the objects which were exported to be redefined.

The first IMPORT command causes the base cluster to be imported into the master catalog.

1. The INFILE parameter names the DD statement that describes and allocates the fourth file on the magnetic tape reel containing the portable data set. (Note: Either INFILE or INDATASET is required.)
2. The OUTDATASET parameter identifies the data set being imported. (Note: Either OUTFILE or OUTDATASET is required.) Access Method Services dynamically allocates the data set after it has been defined; this will also provide access to the catalog recovery area.
3. The CATALOG parameter is required, since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

The second IMPORT command causes an alternate index to be imported into the master catalog.

1. The INFILE parameter names the DD statement that describes and allocates the third file on the magnetic tape reel containing the portable data set.
2. The OUTDATASET parameter identifies the data set being imported. Access Method Services dynamically allocates the alternate index after it has been defined; this will also provide access to the catalog recovery area.
3. The CATALOG parameter is required, since the master catalog is password protected. It specifies the name of the master catalog and its update password which is required to define into a protected catalog.

The third IMPORT command causes the previously-exported entry-sequenced VSAM data set to be imported. The data set is imported into the catalog from which it was exported. The portable data set (that is, the copy being imported) replaces the copy that exists in user catalog D27UCAT2. The IMPORT process searches user catalog D27UCAT2 (located through its alias 'D50') for the entry with an entryname of D50.EXAMPLE.ESDS1. It deletes that entry, then redefines a cluster entry using the catalog information obtained from the portable data set. Because the data set was exported with the TEMPORARY attribute (see the previous example), the IMPORT command doesn't need to supply volume information.

Note: The SOURCE DD statement specifies a recordsize (LRECL=3004) because the largest record in the portable data set is 3000 bytes (that is, LRECL = (largest-record size) + 4). Otherwise, the default record size, blocksize minus 4, would be erroneously used by the EXPORT command. In the two previous steps, the default record size was used.

1. The INFILE parameter names the DD statement that describes and allocates the third file on a magnetic tape reel. This is the portable data set to be imported.
2. The OUTFILE parameter names the DD statement that describes and allocates the data set to be imported. Note that since the data set was exported with the TEMPORARY attribute, it exists in the catalog D27UCAT2 at step allocation time.

3. The CATALOG parameter is required, since the catalog is protected. It names the catalog that is to contain the imported data set. The catalog's master password is supplied, and allows the IMPORT process to delete the existing entry in the catalog and redefine a new entry for the entry-sequenced VSAM data set.

The fourth IMPORT command causes the previously-exported key-sequenced VSAM data set, D40.EXAMPLE.KSDS1, to be imported from its copy (that is, the first portable data set on the magnetic tape reel) to a different user catalog than it was exported from, D27UCAT2. The IMPORT command renames the data set and each of its components, as specified with the NEWNAME parameters, and specifies that the data set is to reside on a different volume than that from which it was exported. The high-level qualifier (D50) of the new component's name is the alias name of the user catalog, D27UCAT2.

Note: The SOURCE DD statement describes the portable data set created in STEP1 of the previous example. The blocksize parameter is included (even though it need not be, because the tape has a standard label and the information is contained in the data set header label) to illustrate the fact that the information specified when the data set is imported is required to be the same as was specified when the data set was exported. The LRECL parameter is not required, because the maximum record size is 475 bytes and the default (blocksize minus 4) is adequate. However, by specifying a record size, the default is overridden and virtual storage is more efficiently used. When record size is specified, it is the largest-record size + 4.

1. The INFILE parameter names the DD statement that describes and allocates the portable data set containing the to-be-imported VSAM data set.
2. The OUTDATASET parameter identifies the renamed data set. Access Method Services dynamically allocates the data set after it has been defined.
3. The CATALOG parameter is required since the catalog is protected. It names the catalog that is to contain the imported data set's entry. The catalog's update password is supplied and allows the data set to be imported into the catalog.
4. The OBJECTS parameter identifies the volume that is to contain the imported data set and specifies new names for each of the data set's components. The OBJECTS parameter identifies each entry of the imported data set with its original entryname, then specifies information that is to replace the information found in the portable data set's imported catalog entries.

EXAMPLE 11: DELETING A VSAM CATALOG AND ITS CATALOGED OBJECTS

This example deletes all entries defined in the previous examples and deletes the user catalogs which were defined. The result of this job is the removal of all data sets and catalogs defined in this set of examples.


```

//DELETE JOB ...
//STEP1 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

/* DELETE THE NONVSAM DATA SET EXAMPLE.NONVSAM1 */
DELETE EXAMPLE.NONVSAM1 -
      NOSCRATCH -
      NONVSAM -
      CATALOG(D27UCAT1/UCATUPPW)

/* DELETE VSAM DATA SETS CATALOGED IN */
/* USER CATALOGS D27UCAT1 AND D27UCAT2 */
DELETE (D50.EXAMPLE.ESDS1/ESD1PWMR -
      D50.EXAMNEW.KSDS1/KSD1PSWD -
      D40.MYDATA) -
      PURGE -
      CLUSTER

/* DELETE THE ENTRY-SEQUENCED VSAM */
/* DATA SET EXAMPLE.ESDS2 */
DELETE EXAMPLE.ESDS2/ESD2MRPW -
      PURGE -
      CATALOG(D27UCAT2)

/* DELETE THE KEY-SEQUENCED VSAM */
/* CLUSTER EXAMPLE.KSDS2 */
DELETE EXAMPLE.KSDS2 -
      CATALOG(AMASTCAT/MCATMRPW)

/* DELETE THE RELATIVE-RECORD VSAM */
/* CLUSTER EXAMPLE.RRDS1 */
DELETE EXAMPLE.RRDS1 -
      CLUSTER

/* DELETE VSAM DATA SPACES FROM VOLUME */
/* VSER04 AND, IF THE VOLUME DOESN'T */
/* CONTAIN VSAM OBJECTS, DELETE THE */
/* VOLUME'S ENTRY IN AMASTCAT */
DELETE VSER04 -
      SPACE -
      CATALOG(AMASTCAT/MCATUPPW)

/* DELETE THE VSAM USER CATALOG D27UCAT1 */
DELETE D27UCAT1/UCATMRPW -
      USERCATALOG -
      PURGE

/* DELETE THE VSAM USER CATALOG D27UCAT2 */
DELETE D27UCAT2/UCATMRPW -
      USERCATALOG -
      PURGE
/*

```

Explanation of Commands

The first DELETE command deletes the nonVSAM data set named EXAMPLE.NONVSAM1.

1. The name of the nonVSAM data set is required.
2. The NOSCRATCH parameter specifies that the VTOC entry of the object being deleted is not to be removed.

3. The NONVSAM parameter ensures that the entry being deleted is a nonVSAM data set.
4. The CATALOG parameter is required in this example to identify the catalog that describes the data set to be deleted. The catalog's update password is required to delete a nonVSAM data set.

The second DELETE command deletes VSAM data sets cataloged in two user catalogs, D27UCAT1 and D27UCAT2.

1. The names of the VSAM data sets are specified; in addition, the master passwords of two protected data sets are required. Since the high-level qualifier of the three data set names is the alias of the catalog that describes the data sets, the CATALOG parameter is not required to direct the catalog search.
2. The PURGE parameter causes the entries to be deleted without regard for the retention period.
3. The CLUSTER parameter ensures that the catalog object being deleted is a VSAM data set.

The third DELETE command deletes the entry-sequenced VSAM data set EXAMPLE.ESDS2 from the user catalog, D27UCAT2.

1. The name of the entry-sequenced data set, EXAMPLE.ESD2, and its master password, ESD2MRPW, are required to identify the object to be deleted.
2. The PURGE parameter causes the entry to be deleted without regard to its retention period.
3. The CATALOG parameter is required in this example to identify the catalog that describes the data set to be deleted.

The fourth DELETE command deletes the key-sequenced VSAM data set EXAMPLE.KSDS2 from the master catalog. Because this data set is a base cluster, its related alternate index, EXAMPLE.AIX, and its path, EXAMPLE.PATH, are also automatically deleted from the master catalog.

1. The name of the key-sequenced VSAM data set, EXAMPLE.KSDS2, is required to identify the object to be deleted.
2. The CATALOG parameter is required, because the delete process will result in more than one protected object being deleted. It names the catalog that contains the catalog entry of each object to be deleted and supplies the catalog's master password, which allows each object in the catalog to be deleted.

The fifth DELETE command deletes the relative-record VSAM data set EXAMPLE.RRDS1 from the master catalog.

1. The name of the relative-record VSAM data set, EXAMPLE.RRDS1, is required to identify the object to be deleted. Because the data set was created without passwords, no password is required to delete it.
2. The CLUSTER parameter specifies that a cluster is being deleted.

The sixth DELETE command deletes each empty VSAM data space from volume VSER04. Because the volume no longer contains any VSAM objects, the DELETE command also removes the volume's entry from the master catalog and removes the master catalog's ownership of the volume. Access Method Services will dynamically allocate the volume, since the VTOC must be accessed. The volume should be mounted permanently RESIDENT or RESERVED to assure successful dynamic allocation.

Note: If VSAM data sets remain on a volume, the DELETE SPACE command only frees unused VSAM space on the volume.

1. The name (that is, volume serial number) of the volume, VSER04, is required to identify the object to be deleted.
2. The SPACE parameter is required to specify that all data spaces on the volume are to be deleted and, if the volume does not contain VSAM objects, that the volume's entry is to be deleted from the catalog.
3. The CATALOG parameter is required and names the catalog that owns the volume. The catalog's update password, MCATUPPW, allows the DELETE command to delete a volume entry or VSAM data space.

The seventh DELETE command deletes the VSAM user catalog D27UCAT1.

1. The name of the catalog and its master password are required to identify the object to be deleted.
2. The USERCATALOG parameter is required to specify that a user catalog is being deleted.
3. The PURGE parameter causes the entry to be deleted without regard to its retention period.

The eighth DELETE command deletes the VSAM user catalog D27UCAT2.

1. The name of the user catalog and its master password are required.
2. The USERCATALOG parameter is required to specify that the object being deleted is a user catalog.
3. The PURGE parameter causes the entry to be deleted without regard for the retention period.

APPENDIX B. INTERPRETING LISTCAT OUTPUT LISTINGS

The various LISTCAT command options allow you to select the LISTCAT output that gives you the information you want. This appendix provides information on the structure of LISTCAT output when you specify certain options. It also lists and describes fields that can be printed for each type of catalog entry.

Each listed entry is identified by its type (that is, cluster, nonVSAM, data, etc.) and by its entryname. Entries are listed in alphabetic order of the entrynames, except when the ENTRIES parameter is used. The entries are then listed in the order in which they are specified in the ENTRIES parameter.

An entry which has associated entries is immediately followed by the listing of each associated entry, unless type options (CLUSTER, DATA, SPACE, etc.) have been specified or a generic entryname list was specified which excludes the associated entry. That is, a cluster's data component (and, if the cluster is key-sequenced, its index component) is listed immediately following the cluster.

This appendix is organized in three parts:

- "LISTCAT Output Keywords," which lists all field names that can be listed for each type of entry.
- "Description of Keyword Fields," which describes each field name within a group of related field names.
- "Examples of LISTCAT Output Listings," which describes and illustrates the LISTCAT output that results when various LISTCAT options are specified.

LISTCAT OUTPUT KEYWORDS

This section of the appendix lists the field names associated with each type of catalog entry. Each field name is followed by an abbreviation that points to a group of related-field descriptions in the next section. Keywords are listed in alphabetic order, not in the order of appearance in the LISTCAT output.

The group names and abbreviations are:

Abbreviations	Group Names
ALC	Allocation Group
ASN	Associations Group
ATT	Attributes Group
DSP	Data Space Group
GDG	Generation Data Group Base Entry, Special Fields for
NVS	NonVSAM Entry, Special Field for
HIS	History Group

Abbreviations	Group Names
PRT	Protection Group
STA	Statistics Group
VLS	Volumes Group
VOL	Volume Entry, Special Fields for

Alias Entry Keywords

ASSOCIATIONS (ASN)
 entryname (HIS)
 HISTORY (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)

Alternate-Index Entry Keywords

ASSOCIATIONS (ASN)
 ATTEMPTS (PRT)
 ATTRIBUTES (ATT)
 CLUSTER (ASN)
 CODE (PRT)
 CONTROLPW (PRT)
 DATA (ASN)
 entryname (HIS)
 HISTORY (HIS)
 CREATION (HIS)
 EXPIRATION (HIS)
 OWNER-IDENT (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
 INDEX (ASN)
 MASTERPW (PRT)
 NOUPGRADE (ATT)
 PATH (ASN)
 PROTECTION (PRT)
 RACF (PRT)
 READPW (PRT)
 UPDATEPW (PRT)
 UPGRADE (ATT)
 USAR (PRT)
 USVR (PRT)

Cluster Entry Keywords

AIX (ASN)
 ASSOCIATIONS (ASN)
 ATTEMPTS (PRT)
 CODE (PRT)
 CONTROLPW (PRT)
 DATA (ASN)
 entryname (HIS)
 HISTORY (HIS)
 CREATION (HIS)
 EXPIRATION (HIS)
 OWNER-IDENT (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
 INDEX (ASN)
 MASTERPW (PRT)
 PATH (ASN)
 PROTECTION (PRT)

RACF (PRT)
READPW (PRT)
UPDATEPW (PRT)
USAR (PRT)
USVR (PRT)

Data Entry Keywords

ALLOCATION (ALC)
AIX (ASN)
ASSOCIATIONS (ASN)
ATTEMPTS (PRT)
ATTRIBUTES (ATT)
AVGLRECL (ATT)
AXRKP (ATT)
BIND (ATT)
BUFSPACE (ATT)
BYTES/TRACK (VLS)
CI/CA (ATT)
CISIZE (ATT)
CLUSTER (ASN)
CODE (PRT)
CONTROLPW (PRT)
CYLFAULT (ATT)
DEVTYPE (VLS)
DSTGWAIT (ATT)
entryname (HIS)
ERASE (ATT)
EXCPEXIT (ATT)
EXCPS (STA)
EXTENT-NUMBER (VLS)
EXTENT-TYPE (VLS)
EXTENTS (STA)
EXTENTS (VLS)
 HIGH-CCHH (VLS)
 HIGH-RBA (VLS)
 LOW-CCHH (VLS)
 LOW-RBA (VLS)
 TRACKS (VLS)
FREESPACE-%CI (STA)
FREESPACE-%CA (STA)
FREESPC-BYTES (STA)
HIGH-KEY (VLS)
HI-KEY-RBA (VLS)
HISTORY (HIS)
 CREATION (HIS)
 EXPIRATION (HIS)
 OWNER-IDENT (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
IMBED (ATT)
INDEX (ASN)
INH-UPDATE (ATT)
INDEXED (ATT)
KEYLEN (ATT)
LOW-KEY (VLS)
MASTERPW (PRT)
MAXLRECL (ATT)
MAXRECS (ATT)
NOERASE (ATT)
NOIMBED (ATT)
NONINDEXED (ATT)
NOREPLICAT (ATT)
NOREUSE (ATT)
NONSPANNED (ATT)
NOSWAP (ATT)
NOTRKOVL (ATT)
NONUNIQKEY (ATT)
NOTUSABLE (ATT)
NOWRITECHK (ATT)
NUMBERED (ATT)

ORDERED (ATT)
 PGSPC (ASN)
 PHYRECS/TRK (VLS)
 PHYREC/SIZE (VLS)
 PROTECTION (PRT)
 RACF (PRT)
 HI-ALLOC-RBA (ALC)
 HI-USED-RBA (ALC)
 HI-ALLOC-RBA (VLS)
 HI-USED-RBA (VLS)
 READPW (PRT)
 RECOVERY (ATT)
 REC-DELETED (STA)
 REC-INSERTED (STA)
 REC-RETRIEVED (STA)
 REC-TOTAL (STA)
 REC-UPDATED (STA)
 RECORDS/CI (ATT)
 REPLICATE (ATT)
 RKP (ATT)
 REUSE (ATT)
 RECVABLE (ATT)
 SHROPTNS (ATT)
 SPACE-PRI (ALC)
 SPACE-SEC (ALC)
 SPACE-TYPE (ALC)
 SPEED (ATT)
 SPLITS-CA (STA)
 SPLITS-CI (STA)
 SPANNED (ATT)
 STATISTICS (STA)
 SUBALLOC (ATT)
 SWAP (ATT)
 SYSTEM-TIMESTAMP (STA)
 TEMP-EXP (ATT)
 TRACKS/CA (VLS)
 TRKOVFL (ATT)
 UNORDERED (ATT)
 UNIQUE (ATT)
 UNIQUEKEY (ATT)
 UPDATEPW (PRT)
 USAR (PRT)
 USVR (PRT)
 VOLFLAG (VLS)
 VOLSER (VLS)
 VOLUMES (VLS)
 WRITECHECK (ATT)

Index Entry Keywords

ALLOCATION (ALC)
 AIX (ASN)
 ASSOCIATIONS (ASN)
 ATTEMPTS (PRT)
 ATTRIBUTES (ATT)
 AVGLRECL (ATT)
 BIND (ATT)
 BUFSPACE (ATT)
 CI/CA (ATT)
 CISIZE (ATT)
 CLUSTER (ASN)
 CODE (PRT)
 CONTROLPW (PRT)
 CYLFAULT (ATT)
 DEVTYPE (VLS)
 DSTGWAIT (ATT)
 entryname (HIS)
 ERASE (ATT)
 EXCPEXIT (ATT)
 EXTENTS (STA)
 EXTENT-NUMBER (VLS)
 EXTENT-TYPE (VLS)
 EXTENTS (VLS)

HIGH-CCHH (VLS)
 HIGH-RBA (VLS)
 LOW-CCHH (VLS)
 LOW-RBA (VLS)
 TRACKS (VLS)
 FREESPACE-%CI (STA)
 FREESPACE-%CA (STA)
 FREESPC-BYTES (STA)
 HIGH-KEY (VLS)
 HISTORY (HIS)
 CREATION (HIS)
 EXPIRATION (HIS)
 OWNER-IDENT (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
 IMBED (ATT)
 INDEX (STA)
 ENTRIES/SECT (STA)
 HI-LEVEL-RBA (STA)
 LEVELS (STA)
 SEQ-SET-RBA (STA)
 INH-UPDATE (ATT)
 KEYLEN (ATT)
 LOW-KEY (VLS)
 MASTERPW (PRT)
 MAXLRECL (ATT)
 NOERASE (ATT)
 NOIMBED (ATT)
 NOREPLICAT (ATT)
 NOREUSE (ATT)
 EXCPS (STA)
 EXTENTS (STA)
 NOTUSABLE (ATT)
 NOWRITECHK (ATT)
 ORDERED (ATT)
 PHYRECS/TRK (VLS)
 PHYREC-SIZE (VLS)
 PROTECTION (PRT)
 RACF (PRT)
 HI-ALLOC-RBA (ALC)
 HI-USED-RBA (ALC)
 HI-ALLOC-RBA (VLS)
 HI-USED-RBA (VLS)
 RECOVERY (ATT)
 READPW (PRT)
 REC-DELETED (STA)
 REC-INSERTED (STA)
 REC-RETRIEVED (STA)
 REC-TOTAL (STA)
 REC-UPDATED (STA)
 REPLICATE (ATT)
 RKP (ATT)
 REUSE (ATT)
 SHROPTNS (ATT)
 SPACE-PRI (ALC)
 SPACE-SEC (ALC)
 SPACE-TYPE (ALC)
 SPEED (ATT)
 SPLITS-CA (STA)
 SPLITS-CI (STA)
 STATISTICS (STA)
 SUBALLOC (ATT)
 SYSTEM-TIMESTAMP (STA)
 TEMP-EXP (ATT)
 TRACKS/CA (VLS)
 UNORDERED (ATT)
 UNIQUE (ATT)
 UPDATEPW (PRT)
 USAR (PRT)
 USVR (PRT)
 VOLFLAG (VLS)
 VOLSER (VLS)

VOLUME (VLS)
WRITECHECK (ATT)

Generation Data Group Base Entry Keywords

ASSOCIATIONS (ASN)
ATTRIBUTES (GDG)
EMPTY (GDG)
LIMIT (GDG)
NOEMPTY (GDG)
NOSCRATCH (GDG)
SCRATCH (GDG)
entryname (HIS)
HISTORY (HIS)
CREATION (HIS)
EXPIRATION (HIS)
OWNER-IDENT (HIS)
RCVY-CI (HIS)
RCVY-DEVT (HIS)
RCVY-VOL (HIS)
RELEASE (HIS)
NONVSAM (ASN)

NonVSAM Entry Keywords

ALIAS (ASN)
ASSOCIATIONS (ASN)
DEVTYPE(VLS)
entryname (HIS)
FSEQN (NVS)
HISTORY (HIS)
CREATION (HIS)
EXPIRATION (HIS)
OWNER-IDENT (HIS)
RCVY-CI (HIS)
RCVY-DEVT (HIS)
RCVY-VOL (HIS)
RELEASE (HIS)
VOLSER(VLS)

Page Space Entry Keywords

ASSOCIATIONS (ASN)
ATTEMPTS (PRT)
CODE (PRT)
CONTROLPW (PRT)
DATA (ASN)
entryname (HIS)
HISTORY (HIS)
CREATION (HIS)
EXPIRATION (HIS)
OWNER-IDENT (HIS)
RCVY-CI (HIS)
RCVY-DEVT (HIS)
RCVY-VOL (HIS)
RELEASE (HIS)
INDEX (ASN)
MASTERPW (PRT)
PROTECTION (PRT)
RACF (PRT)
READPW (PRT)
UPDATEPW (PRT)
USAR (PRT)
USVR (PRT)

Path Entry Keywords

AIX (ASN)
ASSOCIATIONS (ASN)
ATTEMPTS (PRT)
ATTRIBUTES (ATT)
CLUSTER (ASN)
CODE (PRT)
CONTROLPW (PRT)
DATA (ASN)
entryname (HIS)
HISTORY (HIS)
 CREATION (HIS)
 EXPIRATION (HIS)
 OWNER-IDENT (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
INDEX (ASN)
MASTERPW (PRT)
NOUPDATE (ATT)
PROTECTION (PRT)
RACF (PRT)
READPW (PRT)
UPDATE (ATT)
UPDATEPW (PRT)
USAR (PRT)
USVR (PRT)

User Catalog Entry Keywords

ALIAS (ASN)
ASSOCIATIONS (ASN)
DEVTYPE(VLS)
entryname (HIS)
HISTORY (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
VOLFLAG (VLS)
VOLSER (VLS)

Volume Entry Keywords

ATTRIBUTES (DSP)
 AUTOMATIC (DSP)
 EXPLICIT (DSP)
 MASTERCAT (DSP)
 SUBALLOC (DSP)
 UNIQUE (DSP)
 USERCAT (DSP)
BYTES/TRK (VOL)
CHARACTERISTICS (VOL)
CYLS/VOL (VOL)
DATASET-DIRECTORY (DSP)
 ATTRIBUTES (DSP)
 DSN (DSP)
 EXTENTS (DSP)
DATASETS (DSP)
DATASETS-ON-VOL (VOL)
DATASPACE (DSP)
DATASPCS-ON-VOL (VOL)
DEVTYPE (VOL)
EXTENT-DESCRIPTOR (DSP)
 BEG-CCHH (DSP)
 SPACE-MAP (DSP)
 TRACKS-TOTAL (DSP)
 TRACKS-USED (DSP)
EXTENTS (DSP)
FORMAT-1-DSCB (DSP)

CCHHR (DSP)
 TIMESTAMP (DSP)
 MAX-PHYREC-SZ (VOL)
 MAX-EXT/ALLOC (VOL)
 HISTORY (HIS)
 RCVY-CI (HIS)
 RCVY-DEVT (HIS)
 RCVY-VOL (HIS)
 RELEASE (HIS)
 SEC ALLOC (DSP)
 TRKS/CYL (VOL)
 TYPE (DSP)
 CYLINDER (DSP)
 TRACK (DSP)
 volume serial number (HIS)
 VOLUME-TIMESTAMP (VOL)

DESCRIPTION OF KEYWORD FIELDS

This section of the appendix contains a description of each field name. The field names are in the following groups of related information:

Abbreviations	Group Names
ALC	Allocation Group
ASN	Associations Group
ATT	Attributes Group
DSP	Data Space Group
GDG	Generation Data Group Base Entry, Special Fields for
NVS	NonVSAM Entry, Special Field for
HIS	History Group
PRT	Protection Group
STA	Statistics Group
VLS	Volumes Group
VOL	Volume Entry, Special Fields for

Groups are in alphabetic order. Field names within each group are in alphabetic order, not the order of appearance in the listed entry.

ALC: ALLOCATION GROUP

The fields in this group describe the space allocated to the data or index component defined by the entry.

HI-ALLOC-RBA—The highest RBA (plus 1) available within allocated space to store data.

HI-USED-RBA—The highest RBA (plus 1) within allocated space that actually contains data.

SPACE-PRI—Gives the number of units (indicated under TYPE) of space allocated to the data or index component when the cluster (that is, its data or index component) was defined.

This amount of space is to be allocated whenever a data component (or key range within it, and its associated sequence set, if IMBED is an attribute of the cluster) is extended onto a candidate volume.

SPACE-SEC—Gives the number of units (indicated under TYPE) of space to be allocated whenever a data set (or key range within it) is extended on the same volume.

SPACE-TYPE—Indicates the unit of space allocation:

CYLINDER—Cylinders

TRACK—Tracks

ASN: ASSOCIATIONS GROUP

This group lists the type (cluster, data, etc.) and entry names of the objects associated with the present entry. A cluster or alternate index entry will indicate its associated path entries and data and index (if a key-sequenced data set) entries. Similarly, an index or data entry will indicate its associated cluster or the alternate index of which it is a component.

- An alias entry points to:
 - Its associated nonVSAM data set entry.
 - A user catalog entry. (All alias entries for a nonVSAM data set entry are chained together, as are alias entries for a user catalog entry.)
- An alternate-index entry points to:
 - Its associated data and index entries.
 - Its base cluster's cluster entry.
 - Each associated path entry.
- An alternate-index's data entry points to:
 - Its associated alternate-index entry.
- An alternate-index's index entry points to:
 - Its associated alternate-index entry.
- A cluster entry points to:
 - Its associated data entry.
 - Each associated path entry.
 - For a key-sequenced cluster, its associated index entry.
 - For a cluster with alternate indexes, each associated alternate index entry.
- A cluster's data entry points to:
 - Its associated cluster entry.
- A cluster's index entry points to:
 - Its associated cluster entry.
- A generation data group base entry points to:
 - Its associated nonVSAM data set entries.
- A nonVSAM data set entry points to:

- Its associated alias entry.
- Its associated generation data group (for a G0000V00 nonVSAM).
- A page space entry points to:
 - Its associated data entry. (The page space is cataloged as an entry-sequenced cluster with a cluster entry and an associated data entry.)
- A path entry (that establishes the connection between a base cluster and an alternate index) points to:
 - Its associated alternate index entry, and the alternate index's associated data and index entries.
 - The data entry of its associated base cluster.
 - For a key-sequenced base cluster, the index entry of its associated base cluster.
- Path entry (that is an alias for a cluster entry) points to:
 - Its associated base cluster entry.
 - The data entry of its associated base cluster.
 - For a key-sequenced cluster, the index entry of its associated base cluster.
- A user-catalog entry points to:
 - Its associated alias entry.

AIX—Identifies an alternate-index entry.

ALIAS—Identifies an alias entry.

CLUSTER—Identifies a cluster entry.

DATA—Identifies a data entry.

GDG—Identifies a generation data group (GDG) base entry.

INDEX—Identifies an index entry.

NONVSAM—Identifies a nonVSAM data set entry.

PGSPC—Identifies a page space entry.

PATH—Identifies a path entry.

UCAT—Identifies a user-catalog entry.

ATT: ATTRIBUTES GROUP

The fields in this group describe the miscellaneous attributes of the entry. See the DEFINE command for further discussion of most of these attributes.

AVGLRECL—The average length of data records. AVGLRECL equals MAXLRECL when the records are fixed-length.

AXRKP—Indicates, for an alternate index, the offset, from the beginning of the base-cluster's data record, at which the alternate-key field begins.

BIND—The component is staged from mass storage to a direct-access storage staging drive when it is opened, and it is retained (bound) in direct-access storage until it is closed.

BUFSPACE—The minimum buffer space in virtual storage to be provided by a processing program.

CI/CA—The number of control intervals per control area.

CISIZE—The size of a control interval.

CYLFAULT—The component is not staged from mass storage to a direct-access storage staging drive when it is opened, but data from it is staged as the data is needed.

DSTGWAIT—The component is destaged from direct-access storage to mass storage before VSAM returns control to the program that is closing it.

ERASE—Records are to be erased (set to binary zeros) when deleted.

EXCPEXIT—The name of the object's exception exit routine.

IMBED—The sequence-set index record is stored along with its associated data control area.

INH-UPDATE—The data component cannot be updated. Either the data component was exported with INHIBITSOURCE specified, or its entry was modified by way of ALTER, with INHIBIT specified.

INDEXED—The data component has an index—it is key-sequenced.

KEYLEN—The length of the key field in a data record.

MAXLRECL—The maximum length of data or index records. AVGLRECL equals MAXLRECL when the records are fixed-length.

MAXRECS—Identifies the highest possible valid relative-record number, for a relative-record data set. This value is calculated as follows:

$$2^{32}/\text{CISIZE} \times \text{number of records slots per control interval.}$$

NOERASE—Records are not to be erased (set to binary 0s) when deleted.

NOIMBED—The sequence-set index record is not stored along with its associated data control area.

NONINDEXED—The data component has no index—it is entry-sequenced.

NONSPANNED—Data records cannot span control intervals.

NONUNIQKEY—Indicates, for an alternate index, that more than one data record in the base cluster can contain the same alternate-key value.

NOREPLICAT—Index records are not replicated.

NOREUSE—The data set cannot be reused.

NOSNAP—The page space is a conventional page space and cannot be used as a high speed swap data set.

NOTRKOVL—The physical blocks of a page space data set cannot span a track boundary.

NOUPDATE—When the path is opened for processing, its associated base cluster is opened but the base cluster's upgrade set is not opened.

NOUPGRADE—The alternate index is not upgraded unless it is opened and being used to access the base cluster's data records.

NOTUSABLE—The entry is not usable because (1) the catalog could not be correctly recovered by RESETCAT, or (2) a DELETE SPACE FORCE was issued for a volume(s) in the entry's volume list.

NOWRITECHK—Write operations are not checked for correctness.

NUMBERED—The cluster is a relative-record data set.

ORDERED—Volumes are used for space allocation in the order they were specified when the cluster was defined.

RECOVERY—A temporary CLOSE is issued as each control area of the data set is loaded, so the whole data set won't have to be reloaded if a serious error occurs during loading.

RECORDS/CI—Specifies the number of records, or slots, in each control interval of a relative-record data set.

RECVABLE—This attribute is set in the data entry of a recoverable catalog, and each of the catalog's volumes contains a catalog recovery area.

REPLICATE—Index records are replicated (that is, each is duplicated around a track of the index's direct-access device).

REUSE—The data set can be reused (that is, its contents are temporary and its high-used RBA can be reset to 0 when it is opened).

RKP—The relative key position—the displacement from the beginning of a data record to its key field.

SHROPTNS—(n,m) The numbers n and m identify the types of sharing permitted. See SHAREOPTIONS in the DEFINE CLUSTER section for more details.

SPANNED—Data records can be longer than control-interval length, and can cross, or span, control-interval boundaries.

SPEED—CLOSE is not issued until the data set has been loaded.

SUBALLOC—More than one VSAM cluster can share the data space. A VSAM catalog might also occupy the data space.

SWAP—The page space is a high speed swap data set used by Auxiliary Storage Management during a swap operation to store and retrieve the set of LSQA pages owned by an address space.

TEMP-EXP—The data component was temporarily exported.

TRKOVFL—The physical blocks of a page space data set can span a track boundary.

UNIQUE—Only one VSAM cluster or catalog can occupy the data space—the cluster or catalog is unique.

UNIQUEKEY—Indicates, for an alternate index, that the alternate-key value identifies one, and only one, data record in the base cluster.

UNORDERED—Volumes specified when the cluster was defined can be used for space allocation in any order.

UPDATE—When the path is opened, the upgrade set's alternate indexes (associated with the path's base cluster) are also opened and are updated when the base cluster's contents change.

UPGRADE—When the alternate index's base cluster is opened, the alternate index is also opened and will be updated to reflect any changes to the base cluster's contents.

WRITECHECK—Write operations are checked for correctness.

DSP: DATA SPACE GROUP

The fields in this group are included by LISTCAT, as part of a volume entry, for each data space on the volume. If a volume contains no data spaces, it is a candidate volume.

ATTRIBUTES—Describes the attributes of the data space.

AUTOMATIC—The data space was created automatically by a secondary allocation operation. If a VSAM cluster must be given additional space on a volume and all existing data spaces are full, an existing data space will be extended, or a new one allocated on the volume, using catalog DADSM. A new data space so allocated is known as an implicit or automatic data space.

EXPLICIT—The data space was created explicitly by a DEFINE MASTERCATALOG or USERCATALOG command, or a DEFINE ALTERNATEINDEX or DEFINE CLUSTER command with the UNIQUE attribute specified or by a DEFINE SPACE command.

MASTERCAT—The data space contains the master catalog.

SUBALLOC—The data space might contain several VSAM clusters.

USERCAT—The data space contains a user catalog.

UNIQUE—The data space contains a single (unique) VSAM cluster or catalog.

DATASET DIRECTORY—Lists the VSAM data sets that can be stored (see CANDIDATE below) or actually are stored (in whole or in part) in the data space.

ATTRIBUTES—Describes the relation between the named data set and the data space.

CANDIDATE—The volume is a candidate for storing the data set

(NULL)—The data set is stored (in whole or in part) in the data space

DSN—The data-set name of the object that can be stored or is stored on the volume.

EXTENTS—The number of data set extents for the data set within the data space.

DATASETS—The number of VSAM data sets stored (in whole or in part) in the data space. (The number includes data sets for which the volume is a candidate.)

EXTENT-DESCRIPTOR—Describes the data space extent.

BEG-CCHH—The device address (that is, CC = cylinder and HH = track) of the extent.

SPACE-MAP—A hexadecimal number that tells what tracks are used and what tracks are free in the extent. The number consists of one or more RLCs (run length codes). The first RLC gives the number of contiguous used tracks, starting at the beginning of the extent; if all of the tracks in the extent are used, there is only one RLC. The second RLC gives the number of contiguous free tracks, beyond the used tracks. A third RLC gives used tracks again, a fourth gives free tracks, and so on.

A 1-byte RLC gives the number of tracks less than 250; a 3-byte RLC gives the number of tracks equal to or greater than 250. That is, if the first byte of an RLC is X'F9' (249) or less, it is the only byte of the RLC and gives the number of tracks. If the first byte of an RLC is X'FA' (250)

or more, the byte is followed by two more bytes that give the number of tracks (the first byte means nothing more than to look at the next two bytes).

TRACKS-TOTAL—Allocated to the data space altogether.

TRACKS-USED—Used to store data.

EXTENTS—The number of data space extents in the data space.

FORMAT-1-DSCB—Identifies the format-1 DSCB that describes the data space.

CCHHR—The device address (that is, CC = cylinder, HH = track, and R = record number) of the DSCB in the VTOC.

TIMESTAMP—The time the data space was allocated (System/370 time-of-day clock value). The DSCB contains the timestamp as part of the VSAM-generated name.

SEC-ALLOC—Gives the number of units (indicated under TYPE) of space to be allocated whenever the data space is extended.

TYPE—Indicates the unit of space allocation:

CYLINDERS—Cylinders

TRACK—Tracks

GDG: GENERATION DATA GROUP BASE ENTRY, SPECIAL FIELDS FOR

The special fields for a generation data group base entry describe attributes of the generation data group.

ATTRIBUTES

This field includes the following fields:

EMPTY

All generation data sets in the generation data group will be uncataloged when the maximum number (given under LIMIT) is reached and one more data set is to be added to the group.

LIMIT

The maximum number of generation data sets allowed in the generation data group.

NOEMPTY

Only the oldest generation data set in the generation data group will be uncataloged when the maximum number (given under LIMIT) is reached and one more data set is to be added to the group.

NOSCRATCH

Generation data sets are not to be scratched (see SCRATCH below) when uncataloged.

SCRATCH

Generation data sets are to be scratched (that is, the DSCB that describes each one is removed from the VTOC of the volume on which it resides) when uncataloged.

NVS: NONVSAM ENTRY, SPECIAL FIELD FOR

The special field for a nonVSAM data set describes a nonVSAM data set stored on magnetic tape.

FSEQN—The sequence number (for the tape volume indicated under the "VOLUMES group" keyword VOLSER) of the file in which the nonVSAM data set is stored.

HIS: HISTORY GROUP

The fields in this group identify the object's owner, identify its catalog recovery volume and release level, and give its creation and expiration dates.

entryname—The name of the cataloged object. The entryname can be specified with the ENTRIES parameter of LISTCAT to identify a catalog entry.

HISTORY—This field includes the following fields:

CREATION—The julian date (YY.DDD) on which the entry was created. This field will always contain 00.000 for a nonVSAM data set entry in an OS CVOL.

EXPIRATION—The julian date (YY.DDD) after which the entry can be deleted without specifying the PURGE parameter in the DELETE command. Julian date 99.999 indicates that PURGE is always required to delete the object. This field will always contain 00.000 for a nonVSAM data set entry in an OS CVOL.

OWNER-IDENT—The identity of the owner of the object described by the entry.

RCVY-CI—If the entry is in a recoverable catalog, this field contains the control interval number in the catalog recovery area where a duplicate of the entry can be found.

RCVY-DEVT—If the entry is in a recoverable catalog, this field identifies the recovery-volume's device type. (See the Device Type Translate Table at the end of this section for a translation of the hexadecimal device type to its equivalent EBCDIC value.)

RCVY-VOL—If the entry is in a recoverable catalog, this field contains the recovery-volume's serial number.

RELEASE—The release of VSAM under which the entry was created (not the same as the release number of OS/VS2):

1 = OS/VS2 Releases 3 and releases previous to Release 3

2 = OS/VS2 Release 3.6 and any later releases

PRT: PROTECTION GROUP

The fields in this group describe how the alternate index, cluster, data component, index component, or path defined by the entry is password protected or RACF protected. NULL or SUPPRESSED might be listed under password protection and YES or NO might be listed under RACF protection.

NULL indicates that the object defined by the entry has no passwords.

SUPP indicates that the master password of neither the catalog nor the entry was specified, so authority to list protection information is not granted.

RACF—Indicates whether or not the entry is protected by the RESOURCE Access Control Facility.

YES—Entry is RACF protected.

NO—Entry is not RACF protected.

ATTEMPTS—Gives the number of times the console operator is allowed to attempt to enter a correct password.

CODE—Gives the code used to tell the console operator what alternate index, catalog, cluster, path, data component, or index component requires him to enter a password. NULL is listed

under CODE if a code is not used—the object requiring the password is identified with its full name.

CONTROLPW—The control-interval password (that is, the password for control-interval access). NULL indicates no control-interval password.

MASTERPW—The master password.

READPW—The read-only password. NULL indicates no read-only password.

UPDATEPW—The update password. NULL indicates no update password.

USAR—The contents (1 to 255 bytes, in character format) of the USAR (user-security-authorization record). This is the information specified in the string subparameter of the AUTH subparameter of the DEFINE command.

USVR—The name of the USVR (user-security-verification routine) which is to be invoked to verify authorization of any access to the entry.

STA: STATISTICS GROUP

The fields in this group give numbers and percentages that tell how much activity has taken place in the processing of a data or index component. The statistics in the catalog are updated when the data set is closed. Therefore, if a failure occurs during CLOSE, the statistics may not be valid. Since a VSAM catalog remains open from job to job, its statistics are not updated.

Statistic values can be any value from a minus to a large number, if a failure occurs before or during close processing. VERIFY does not correct these statistics.

FREESPACE-%CI—Percentage of space to be left free in a control interval for subsequent processing.

FREESPACE-%CA—Percentage of control intervals to be left free in a control area for subsequent processing.

FREESPC-BYTES—Actual number of bytes of free space in the total amount of space allocated to the data or index component. Free space in partially used control intervals is not included in this statistic.

INDEX—This field appears only in an index entry. The fields under it describe activity in the index component.

ENTRIES/SECT—The number of entries in each section of entries in an index record.

HI-LEVEL-RBA—The RBA (relative byte address) of the highest-level index record.

LEVELS—The number of levels of records in the index. The number is 0 if no records have been loaded into the key-sequenced data set to which the index belongs.

SEQ-SET-RBA—The RBA (relative byte address), in decimal, of the first sequence-set record. The sequence set might be separated from the index set by some quantity of RBA space.

The remaining fields in the Statistics group (except for the System Timestamp field), are updated only when the data set is closed.

EXCPS—EXCP (execute channel program—SVC 0) macro instructions issued by VSAM against the data or index component.

EXTENTS—Extents in the data or index component.

REC-DELETED—The number of records that have been deleted from the data or index component. Statistics for records deleted are not maintained when the data set is processed in control interval mode.

REC-INSERTED—For a key-sequenced data set, the number of records that have been inserted into the data component before the last record; records originally loaded and records added to the end are not included in this statistic. For relative-record data sets, it is the number of records inserted into available slots; the number of records originally loaded are included in this statistic. Statistics for records inserted are not maintained when the data set is processed in control interval mode.

REC-RETRIEVED—The number of records that have been retrieved from the data or index component, whether for update or not for update. Statistics for records retrieved are not maintained when the data set is processed in control interval mode.

REC-TOTAL—The total number of records actually in the data or index component. This statistic is not maintained when the data set is processed in control interval mode.

REC-UPDATED—The number of records that have been retrieved for update and rewritten. This value does not reflect those records that were deleted, but a record that is updated and then deleted is counted in the update statistics. Statistics for records updated are not maintained when the data set is processed in control interval mode.

SPLITS-CA—Control-area splits. Half the data records in a control area were written into a new control area and then were deleted from the old control area.

SPLITS-CI—Control-interval splits. Half the data records in a control interval were written into a new control interval and then were deleted from the old control interval.

SYSTEM-TIMESTAMP—The time (System/370 time-of-day clock value) the data or index component was last closed (after being opened for operations that might have changed its contents).

VLS: VOLUMES GROUP

The fields in this group identify the volume(s) on which a data component, index component, user catalog, or nonVSAM data set is stored. It also identifies candidate volume(s) for a data or index component. The fields describe the type of volume and give, for a data or index component, information about the space the object uses on the volume.

- If an entry-sequenced or relative-record cluster's data component has more than one VOLUMES group, each group describes the extents that contain data records for the cluster on a specific volume.
- If a key-sequenced cluster's data component has more than one VOLUMES group, each group describes the extents that contain data records for the cluster, or one of its key ranges, on a specific volume.
- If a key-sequenced cluster's index component has more than one VOLUMES group, each group describes the extents that contain index records for the cluster, or one of its key ranges, on a specific volume. The first VOLUMES group describes the extent that contains the high-level index records (that is, index records in levels above the sequence set level). Each of the next groups describes the extents that contain sequence-set index records for the cluster, or one of its key ranges, on a specific volume. The index

component of a key-sequenced data set with the IMBED attribute will have a minimum of two Volume Groups, one for the imbedded sequence set, and one for the high-level index. The extents for the imbedded sequence set will be the same as those for the data component.

BYTES/TRACK—The number of bytes that VSAM can write on a track (listed for page spaces only).

DEVTYPE—The type of device to which the volume belongs. (See the Device Type Translate Table at the end of this section for a translation of the hexadecimal device type to its equivalent EBCDIC value.)

EXTENT-NUMBER—The number of extents allocated for the data or index component on the volume.

EXTENT-TYPE—The type of extents:

00—The extents are contiguous.

40—The extents are not preformatted.

80—A sequence set occupies a track adjacent to a control area.

EXTENTS—Gives the physical and relative-byte addresses of each extent.

HIGH-CCHH—The device address (that is, CC = cylinder and HH = track) of the end of the extent.

HIGH-RBA—The RBA (relative byte address), in decimal, of the end of the extent.

LOW-CCHH—The device address (that is, CC = cylinder and HH = track) of the beginning of the extent.

LOW-RBA—The RBA (relative byte address), in decimal, of the beginning of the extent.

TRACKS—The number of tracks in the extent, from low to high device addresses.

HIGH-KEY³—For a key-sequenced data set with the KEYRANGE attribute, the highest hexadecimal value allowed on the volume in the key field of a record in the key range. A maximum of 64 bytes can appear in HIGH-KEY.

HI-KEY-RBA³—For a key-sequenced data set, the RBA (relative byte address), in decimal, of the control interval on the volume that contains the highest keyed record in the data set or key range.

LOW-KEY³—For a key-sequenced data set with the KEYRANGE attribute, the lowest hexadecimal value allowed on the volume in the key field of a record in the key range. A maximum of 64 bytes can appear in LOW-KEY.

PHYRECS/TRK—The number of physical records (of the size indicated under PHYRECS-SIZE) that VSAM can write on a track on the volume.

PHYREC-SIZE—The number of bytes that VSAM uses for a physical record in the data or index component.

HI-ALLOC-RBA—The highest RBA (plus 1) available within allocated space to store data component, its key range, the index component, or the sequence set records of a key range.

³ Multiple key ranges might reside on a single volume—the volumes group is repeated for each such key range field.

HI-USED-RBA—The highest RBA (plus 1) within allocated space that actually contains data component, its key range, the index component, or the sequence set records of a key range.

TRACKS/CA—The number of tracks in a control area in the data component. (This value is computed when the entry is defined, and reflects the optimum size of the control area for the given device type and the nature of the entry—whether indexed, nonindexed, or numbered.) For a key-sequenced data set with the imbedded attribute, this value includes the sequence set track.

VOLFLAG—Indicates whether the volume is a candidate volume or the first or a subsequent volume on which data in a given key range is stored.

CANDIDATE—The volume is a candidate for storing the data or index component.

OVERFLOW—The volume is an overflow volume on which data records in a key range are stored. The KEYRANGE begins on another (PRIME) volume.

PRIME—The volume is the first volume on which data records in a key range are stored.

VOLSER—The serial number of the volume.

VOL: VOLUME ENTRY, SPECIAL FIELDS FOR

The special fields for a volume entry describe the characteristics of the space that VSAM uses on the volume.

volume serial number—The name of the cataloged volume entry. The volume serial number can be specified with the ENTRIES parameter of LISTCAT to identify the volume entry.

BYTES/TRK—The number of bytes that VSAM can use on each track on the volume.

CYLS/VOL—The number of cylinders that VSAM can use on the volume, including alternate track cylinders.

DATASETS-ON-VOL—The number of VSAM clusters that reside, in whole or in part, on the volume. Should the number of data spaces be zero, then this is a candidate volume only and this group is omitted.

DATASPCS-ON-VOL—The number of VSAM data spaces on the volume. Should the number of data spaces be zero, then this is a candidate volume only, and the data space group is omitted.

DEVTYPE—The type of device to which the volume belongs. (See the Device Type Translate Table at the end of this section for a translation of the hexadecimal device type to its equivalent EBCDIC value.)

MAX-PHYREC-SZ—The size of the largest physical record that VSAM can write on the volume.

MAX-EXT/ALLOC—The maximum number of extents that can be allocated per each allocation request for a single data set on the volume.

TRKS/CYL—The number of tracks in each cylinder on the volume.

VOLUME-TIMESTAMP—The time (System/370 time-of-day clock value) VSAM last changed the contents of the volume. The format-4 DSCB contains the timestamp at offset 76 (X'4C').

DEVICE TYPE TRANSLATE TABLE

LISTCAT	Code	Device Type
3000	8001	9 TRK TAPE
3040	200A	3340 (35M/70M)
3050	2006	2305-1
3050	2007	2305-2
3050	2009	3330
3050	200B	3350
3050	200D	3330-1
3058	2009	3330V
3080	8001	7 TRK TAPE
30C0	2008	2314/2319
3200	8003	9 TRK, 6250 BPI TAPE
3400	8003	9 TRK, 1600 BPI TAPE

Refer to OS/VS2 System Programming Library: System Debugging Handbook, Volumes 1 and 2 for additional device types.

EXAMPLES OF LISTCAT OUTPUT LISTINGS

This section of the appendix illustrates the kind of output you can get when you specify LISTCAT parameters. It also describes the job control language you can specify and the output messages you get when the LISTCAT procedure executes successfully.

JOB CONTROL LANGUAGE (JCL) FOR LISTCAT JOBS

The job control language (JCL) statements that can be used to list a catalog's entries are:

```
//LISTCAT JOB ...
//STEP1 EXEC PGM=IDCAMS
//STPCAT DD DSN=YOURCAT,DISP=SHR
//OUTDD DD DSN=LISTCAT.OUTPUT,UNIT=2400-3,
// VOL=SER=TAPE10,LABEL=(1,NL),DISP=(NEW,KEEP),
// DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629)
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
LISTCAT -
CATALOG(YOURCAT/PASSWORD) -
OUTFILE(OUTDD) -
...
/*
```

The JOB statement contains user and accounting information required for your installation.

The EXEC statement identifies the program to be executed, IDCAMS (that is, the Access Method Services program).

- STEPCAT DD, which allocates your catalog. The catalog can be allocated dynamically if the STEPCAT and JOBCAT are omitted. If the catalog is dynamically allocated, its volume must be mounted as permanently RESIDENT or RESERVED.
 - OUTDD, which specifies an alternate output file, so that the LISTCAT output can be written onto an auxiliary storage device. The LISTCAT command's OUTFILE parameter points to the OUTDD DD statement. Only the LISTCAT output is written to the alternate output device. JCL statements, system messages, and job statistics are written to the SYSPRINT output device.
- DSN=LISTCAT.OUTPUT specifies the name for the magnetic tape file.

- UNIT=2400-3 and VOL=SER=TAPE10 specifies that the file is to be contained on magnetic tape volume TAPE10.
- LABEL=(1,NL) specifies that this is the first file on a nonlabeled tape. You can also use a standard-labeled tape by specifying LABEL=(1,SL). If subsequent job steps produce additional files of LISTCAT output on the same tape volume, you should increment the file number in each job step's LABEL subparameter (that is, LABEL=(2,NL) for the second job step, LABEL=(3,NL) for the third job step, etc.)
- DISP=(NEW,KEEP) specifies that this is a new tape file and is to be rewound when the job finishes. If a subsequent job step prints the tape, DISP=(NEW,PASS) should be specified. If your job step contains more than one LISTCAT command, DISP=(MOD,KEEP) or DISP=(MOD,PASS) can be used to concatenate all of the LISTCAT output in one sequential file.
- DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629) specifies that the LISTCAT output records are variable-length, blocked 5-to-1, and are preceded by an ANSI print-control character.
- SYSPRINT DD, which is required for each Access Method Services job step. It identifies the output queue, SYSOUT=A, on which all LISTCAT output and system output messages are printed (unless the OUTFILE parameter and its associated DD statement is specified—see OUTDD above).

Note: If you want all output to be written to an auxiliary storage device, replace 'OUTDD' with 'SYSPRINT' in the OUTDD DD statement and omit the SYSPRINT DD SYSOUT=A statement.

- SYSIN DD, which specifies, with an asterisk (*), that the statements that follow are the input data statements. A '/*' terminates the input data statements.

The LISTCAT command parameters shown above are common to the LISTCAT examples that follow. Other LISTCAT parameters are coded with each example and the output that results is illustrated. These two parameters are optional:

CATALOG, which identifies the catalog, YOURCAT, whose entries are to be listed. If the catalog is password protected, its read (or higher level) password, PASSWORD, is required. If the passwords and protection attributes of each entry are to be listed, the catalog's master password is required.

OUTFILE, which points to the OUTDD DD statement. The OUTDD DD statement allocates an alternate output file for the LISTCAT output.

If you want to print the LISTCAT output that is contained on an alternate output file, you can use the IEBGENER program. The following shows the JCL required to print the alternate output file, LISTCAT.OUTPUT, that was allocated previously:

```
//PRINTOUT JOB ...
//STEP1 EXEC PGM=IEBGENER
//SYSUT1 DD DSN=LISTCAT.OUTPUT,UNIT=2400-3,
// VOL=SER=TAPE10,LABEL=(1,NL),DISP=(OLD,KEEP),
// DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629)
//SYSUT2 DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
/*
```


LISTCAT AND ACCESS METHOD SERVICES OUTPUT MESSAGES

When the LISTCAT job completes, Access Method Services provides messages and diagnostic information. If an error occurred, an analysis of the error message can be found in OS/VS Message Library: VS2 System Messages. When your LISTCAT job completes successfully, Access Method Services provides messages that follow the entry listing (see Figure 16):

LISTING FROM CATALOG -- MJKCAT

```
THE NUMBER OF ENTRIES PROCESSED WAS:
AIX -----1
ALIAS -----1
CLUSTER -----2
DATA -----3
GDG -----1
INDEX -----3
NONVSAM -----2
PAGESPACE -----0
PATH -----1
SPACE -----1
USERCATALOG -----0
TOTAL -----15
```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE, MAXIMUM CONDITION CODE WAS 0

Figure 16. Messages that Follow the Entry Listing

The first line identifies the catalog that contained the listed entries.

The next group of lines specify the number of each entry-type, and the total number of entries, that were listed. This statistical information can help you determine the approximate size, in records, of your catalog.

The next line specifies the number of entries that could not be listed because the appropriate password was not specified.

The last two messages indicate that the LISTCAT command (FUNCTION) and the job step (IDCAMS) completed successfully. When LISTCAT is invoked from a TSO terminal, IDC0001I is not printed.

LISTCAT OUTPUT LISTING

When you specify LISTCAT without any parameters, the entryname and type of each entry is listed (see Figure 17). The same listing would result if the NAMES parameter were specified.

You can use this type of listing to list the name of each cataloged object and to determine the number of entries in the catalog. The total number of entries is an approximate size, in records, of your catalog.

```
/* LIST ENTRY NAMES OF THE CATALOG */
LISTCAT -
CATALOG (MJKCAT)
```

LISTING FROM CATALOG -- MJKCAT

```
AIX ----- MJK.ALT.INDEX1
  DATA ----- MJK.ALT.INDEX1.DATA
  INDEX ----- MJK.ALT.INDEX1.INDEX
  PATH ----- MJK.AIX1.PATH
CLUSTER ----- MJK.CLUSTER1
  DATA ----- MJK.CLUSTER.DATA
  INDEX ----- MJK.CLUSTER1.INDEX
GDG BASE ----- MJK.GDG1
  NONVSAM ----- MJK.GDG1.G0001V00
ALIAS ----- MJK.GDG1.ALIAS
NONVSAM ----- MJK.NONVSAM1
CLUSTER ----- MJKCAT
  DATA ----- VSAM.CATALOG.BASE.DATA.RECORD
  INDEX ----- VSAM.CATALOG.BASE.INDEX.RECORD
VOLUME ----- 333001
```

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

```
AIX -----1
ALIAS -----1
CLUSTER -----2
DATA -----3
GDG -----1
INDEX -----3
NONVSAM -----2
PAGESPACE -----0
PATH -----1
SPACE -----1
USERCATALOG -----0
TOTAL -----15
```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 17. An Example of LISTCAT Output When No Parameters are Specified

LISTCAT VOLUME OUTPUT LISTING

When the LISTCAT command is specified with the VOLUME parameter, the volume serial number and device type of each volume that contains part or all of the cataloged object are listed (see Figure 18).

```

/* LIST VOLUMES FOR EACH ENTRY */
LISTCAT -
  VOLUME -
  CATALOG (MJKCAT)

LISTING FROM CATALOG -- MJKCAT

AIX ----- MJK.ALT.INDEX1
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000010'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

DATA ----- MJK.ALT.INDEX1.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000011'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

INDEX ----- MJK.ALT.INDEX1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000012'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

PATH ----- MJK.AIX1.PATH
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000014'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

CLUSTER ----- MJK.CLUSTER1
HISTORY
OWNER-IDENT-----OMNCLUST      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000E'
RELEASE-----2              EXPIRATION-----77.140      RCVY-DEVT----X'30502009'

DATA ----- MJK.CLUSTER1.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000D'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

INDEX ----- MJK.CLUSTER1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000F'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

GDG BASE ----- MJK.GDG1
HISTORY
OWNER-IDENT-----OMNGDG        CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000015'
RELEASE-----2              EXPIRATION-----99.365      RCVY-DEVT----X'30502009'

NONVSAM ---- MJK.GDG1.G0001V00
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000016'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'
VOLSER-----333002          DEVTTYPE-----X'30502009'

```

Figure 18 (Part 1 of 2). An Example of LISTCAT VOLUME Output

```

NONVSAM ---- MJK.NONVSAMI
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000018'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'
CLUSTER ----- MJKCAT
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040
RELEASE-----2              EXPIRATION-----99.999

DATA ----- VSAM.CATALOG.BASE.DATA.RECORD
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2              EXPIRATION-----00.000
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

INDEX ----- VSAM.CATALOG.BASE.INDEX.RECORD
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2              EXPIRATION-----00.000
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

VOLUME ----- 333001
HISTORY
RELEASE-----2              RCVY-VOL-----333001      RCVY-DEVT----X'30502009'      RCVY-CI-----X'000009'
VOLUMES
VOLSER-----333001          DEVTTYPE-----X'30502009'

```

LISTING FROM CATALOG -- MJKCAT

```

THE NUMBER OF ENTRIES PROCESSED WAS:
AIX -----1
ALIAS -----0
CLUSTER -----2
DATA -----3
GDG -----1
INDEX -----3
NONVSAM -----2
PAGESPACE -----0
PATH -----1
SPACE -----1
USERCATALOG -----0
TOTAL -----14

```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 18 (Part 2 of 2). An Example of LISTCAT VOLUME Output

LISTCAT SPACE ALL OUTPUT LISTING

When the LISTCAT command is specified with the SPACE and ALL parameters, all the information for each volume entry in the catalog is listed (see Figure 19). You can use this type of listing to determine how space on each cataloged volume is allocated to VSAM data spaces. You may have to list the volume's table of contents (VTOC) to determine how all of the volume's space is allocated.

```

/* LIST VOLUME INFORMATION FOR THE VOLUMES -
   CONTROLLED BY THE CATALOG */
LISTCAT -
SPACE -
ALL -
CATALOG(MJKCAT)
  
```

LISTING FROM CATALOG -- MJKCAT

```

VOLUME ----- 333001
HISTORY
RELEASE-----2          RCVY-VOL-----333001      RCVY-DEVT----X'30502009'    RCVY-CI-----X'000009'
CHARACTERISTICS
BYTES/TRK-----13165    DEVTYPE-----X'30502009'    MAX-PHYREC-SZ-----13030    DATASETS-ON-VOL-----5
TRKS/CYL-----19       VOLUME-TIMESTAMP:          MAX-EXT/ALLOC-----5       DATASPCS-ON-VOL-----1
CYLS/VOL-----411      X'8A51B6A76850E000'
DATASPACE
DATASETS-----5        FORMAT-1-DSCB:              ATTRIBUTES:
EXTENTS-----1         CCHHR-----X'0000000103'    SUBALLOC
SEC-ALLOC-----40      TIMESTAMP                    EXPLICIT
TYPE-----TRACK        X'8A51B61D97B04000'        USERCAT
EXTENT-DESCRIPTOR:
TRACKS-TOTAL-----160   BEG-CCHH-----X'00000002'    SPACE-MAP-----5947
TRACKS-USED-----89
DATASET-DIRECTORY:
DSN----MJKCAT          ATTRIBUTES----- (NULL)    EXTENTS-----3
DSN----MJK.CLUSTER.DATA  ATTRIBUTES----- (NULL)    EXTENTS-----1
DSN----MJK.CLUSTER1.INDEX  ATTRIBUTES----- (NULL)    EXTENTS-----1
DSN----MJK.ALT.INDEX1.DATA  ATTRIBUTES----- (NULL)    EXTENTS-----1
DSN----MJK.ALT.INDEX1.INDEX  ATTRIBUTES----- (NULL)    EXTENTS-----1
  
```

LISTING FROM CATALOG -- MJKCAT

```

THE NUMBER OF ENTRIES PROCESSED WAS:
AIX -----0
ALIAS -----0
CLUSTER -----0
DATA -----0
GDG -----0
INDEX -----0
NONVSAM -----0
PAGESPACE -----0
PATH -----0
SPACE -----1
USERCATALOG -----0
TOTAL -----1
  
```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 19. An Example of LISTCAT SPACE ALL Output

LISTCAT ALL OUTPUT LISTING

When you specify the LISTCAT command and include the ALL parameter, all the information for each catalog entry is listed (see Figure 20). This example illustrates the LISTCAT output for each type of catalog entry. You can use this type of listing to obtain all cataloged information (except password and security information) about each entry that is listed. When you want to list an entry's passwords, you must provide the catalog's master password (which results in listing the passwords of each password-protected entry) or each entry's master password.

Note: When ENTRIES is specified, you specify only those entrynames that identify catalog entries which are not volume entries. If a volume serial number is specified with the ENTRIES parameter, entrynames of other entry types cannot also be specified. However, if the ENTRIES parameter is not specified and if entry types are not specified (that is, CLUSTER, SPACE, DATA, etc.), all entries in the catalog, including volume entries, are listed.

```

/* LIST ALL CATALOGED INFORMATION FOR EACH ENTRY */
LISTCAT -
  ALL -
  CATALOG (MJKCAT)

LISTING FROM CATALOG -- MJKCAT

AIX ----- MJK.ALT.INDEX1
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000010'
RELEASE-----2      EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
PROTECTION-PSWD---- (NULL)      RACF----- (NO)
ASSOCIATIONS
DATA-----MJK.ALT.INDEX1.DATA
INDEX-----MJK.ALT.INDEX1.INDEX
CLUSTER--MJK.CLUSTER1
PATH-----MJK.AIX1.PATH
ATTRIBUTES
UPGRADE

DATA ----- MJK.ALT.INDEX1.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000011'
RELEASE-----2      EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
PROTECTION-PSWD---- (NULL)      RACF----- (NO)
ASSOCIATIONS
AIX-----MJK.ALT.INDEX1
ATTRIBUTES
KEYLEN-----5      AVGLRECL-----4086      BUFSIZE-----8704      CISIZE-----4096
RKP-----5      MAXLRECL-----32600      EXCPEXIT----- (NULL)      CI/CA-----18
AXRKP-----45
SHROPTS(1,3) RECOVERY      SUBALLOC      NOERASE      INDEXED      NOWRITECHK      NOIMBED      NOREPLICAT
UNORDERED      NOREUSE      SPANNED      NONUNIQKEY
STATISTICS
REC-TOTAL-----4      SPLITS-CI-----0      EXCPS-----7
REC-DELETED-----0      SPLITS-CA-----0      EXTENTS-----1
REC-INSERTED-----0      FREESPACE-%CI-----0      SYSTEM-TIMESTAMP:
REC-UPDATED-----0      FREESPACE-%CA-----0      X'8A51B6C0C87CF000'
REC-RETRIEVED-----0      FREESPC-BYTES-----69632
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----6      HI-ALLOC-RBA-----73728
SPACE-SEC-----6      HI-USED-RBA-----73728
VOLUME
VOLSER-----333001      PHYREC-SIZE-----4096      HI-ALLOC-RBA-----73728      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'      PHYRECS/TRK-----3      HI-USED-RBA-----73728      EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME      TRACKS/CA-----6
EXTENTS:
LOW-CCHH-----X'00040008'      LOW-RBA-----0      TRACKS-----6
HIGH-CCHH-----X'00040000'      HIGH-RBA-----73728

```

Figure 20 (Part 1 of 7). An Example of LISTCAT ALL Output

```

INDEX ----- MJK.ALT.INDEX1.INDEX
HISTORY
OWNER-IDENT----- (NULL)          CREATION-----77.040          RCVY-VOL-----333001          RCVY-CI-----X'000012'
RELEASE-----2          EXPIRATION-----00.000          RCVY-DEVT----X'30502009'
PROTECTION-PSWD---- (NULL)          RACF----- (NO)
ASSOCIATIONS
AIX-----MJK.ALT.INDEX1
ATTRIBUTES
KEYLEN-----5          AVGRECL-----0          BUFSIZE-----0          CISIZE-----512
RKP-----5          MAXRECL-----505          EXCPEXIT----- (NULL)          CI/CA-----20
SHROPTNS(1,3) RECOVERY          SUBALLOC          NOERASE          NOWRITECHK          NOIMBED          NOREPLICAT          UNORDERED
NOREUSE
STATISTICS
REC-TOTAL-----1          SPLITS-CI-----0          EXCPS-----3          INDEX:
REC-DELETED-----0          SPLITS-CA-----0          EXTENTS-----1          LEVELS-----1
REC-INSERTED-----0          FREESPACE-XCI-----0          SYSTEM-TIMESTAMP:          ENTRIES/SECT-----4
REC-UPDATED-----0          FREESPACE-XCA-----0          X'8A51B6C0C87CF000'          SEQ-SET-RBA-----0
REC-RETRIEVED-----0          FREESPC-BYTES-----9728          HI-LEVEL-RBA-----0
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----1          HI-ALLOC-RBA-----10240
SPACE-SEC-----1          HI-USED-RBA-----512
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512          HI-ALLOC-RBA-----10240          EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'          PHYRECS/TRK-----20          HI-USED-RBA-----512          EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----1
EXTENTS:
LOW-CCHH-----X'0004000E'          LOW-RBA-----0          TRACKS-----1
HIGH-CCHH-----X'0004000E'          HIGH-RBA-----10239
PATH ----- MJK.AIX1.PATH
HISTORY
OWNER-IDENT----- (NULL)          CREATION-----77.040          RCVY-VOL-----333001          RCVY-CI-----X'000014'
RELEASE-----2          EXPIRATION-----00.000          RCVY-DEVT----X'30502009'
PROTECTION-PSWD---- (NULL)          RACF----- (NO)
ASSOCIATIONS
AIX-----MJK.ALT.INDEX1
DATA-----MJK.ALT.INDEX1.DATA
INDEX-----MJK.ALT.INDEX1.INDEX
DATA-----MJK.CLUSTER1.DATA
INDEX-----MJK.CLUSTER1.INDEX
ATTRIBUTES
UPDATE
CLUSTER ----- MJK.CLUSTER1
HISTORY
OWNER-IDENT-----OMNCLUST          CREATION-----77.040          RCVY-VOL-----333001          RCVY-CI-----X'00000E'
RELEASE-----2          EXPIRATION-----77.140          RCVY-DEVT----X'30502009'
PROTECTION
MASTERPW-----MASTCL          UPDATEPW-----UPDCL          CODE-----CODECL          RACF----- (NO)
CONTROLPW-----CNTLCL          READPW-----READCL          ATTEMPTS-----3          USVR----- (NULL)
USAR----- (NONE)
ASSOCIATIONS
DATA-----MJK.CLUSTER1.DATA
INDEX-----MJK.CLUSTER1.INDEX
AIX-----MJK.ALT.INDEX1

```

Figure 20 (Part 2 of 7). An Example of LISTCAT ALL Output


```

DATA ----- MJK.CLUSTER.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCYVY-VOL-----333001      RCYVY-CI-----X'00000D'
RELEASE-----2              EXPIRATION-----00.000      RCYVY-DEVT-----X'30502009'
PROTECTION-PSWD----- (NULL)      RACF----- (NO)
ASSOCIATIONS
CLUSTER--MJK.CLUSTER1
ATTRIBUTES
KEYLEN-----5              AVGLRECL-----80          BUFSPACE-----25088          CISIZE-----12288
RKP-----5              MAXLRECL-----100         EXCPEXIT-----EXITCLUS          CI/CA-----2
SHROPTNS(1,3)  RECOVERY      SUBALLOC      NOERASE          INDEXED      NOWRITECHK          NOIMBED      NOREPLICAT
UNORDERED      NOREUSE          NONSPANNED
STATISTICS
REC-TOTAL-----20          SPLITS-CI-----0          EXCPS-----4
REC-DELETED-----0        SPLITS-CA-----0          EXTENTS-----1
REC-INSERTED-----0       FREESPACE-%CI-----15     SYSTEM-TIMESTAMP:
REC-UPDATED-----0        FREESPACE-%CA-----20     X'8A51B6ACD5B92000'
REC-RETRIEVED-----20     FREESPC-BYTES-----12288
ALLOCATION
SPACE-TYPE-----TRACK      HI-ALLOC-RBA-----24576
SPACE-PRI-----2          HI-USED-RBA-----24576
SPACE-SEC-----2
VOLUME
VOLSER-----333001        PHYREC-SIZE-----4096      HI-ALLOC-RBA-----24576      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'  PHYRECS/TRK-----3        HI-USED-RBA-----24576      EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME        TRACKS/CA-----2
EXTENTS:
LOW-CCHH-----X'00040005'  LOW-RBA-----0          TRACKS-----2
HIGH-CCHH-----X'00040006'  HIGH-RBA-----24575
INDEX ----- MJK.CLUSTER1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCYVY-VOL-----333001      RCYVY-CI-----X'00000F'
RELEASE-----2              EXPIRATION-----00.000      RCYVY-DEVT-----X'30502009'
PROTECTION-PSWD----- (NULL)      RACF----- (NO)
ASSOCIATIONS
CLUSTER--MJK.CLUSTER1
ATTRIBUTES
KEYLEN-----5              AVGLRECL-----0          BUFSPACE-----0          CISIZE-----512
RKP-----5              MAXLRECL-----505         EXCPEXIT-----EXITCLUS          CI/CA-----20
SHROPTNS(1,3)  RECOVERY      SUBALLOC      NOERASE          NOWRITECHK      NOIMBED          NOREPLICAT      UNORDERED
NOREUSE
STATISTICS
REC-TOTAL-----1          SPLITS-CI-----0          EXCPS-----4          INDEX:
REC-DELETED-----0        SPLITS-CA-----0          EXTENTS-----1          LEVELS-----1
REC-INSERTED-----0       FREESPACE-%CI-----0       SYSTEM-TIMESTAMP:          ENTRIES/SECT-----1
REC-UPDATED-----0        FREESPACE-%CA-----0       X'8A51B6ACD5B92000'          SEQ-SET-RBA-----0
REC-RETRIEVED-----0     FREESPC-BYTES-----9728     HI-LEVEL-RBA-----0
ALLOCATION
SPACE-TYPE-----TRACK      HI-ALLOC-RBA-----10240
SPACE-PRI-----1          HI-USED-RBA-----512
SPACE-SEC-----1
VOLUME
VOLSER-----333001        PHYREC-SIZE-----512      HI-ALLOC-RBA-----10240      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'  PHYRECS/TRK-----20       HI-USED-RBA-----512        EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME        TRACKS/CA-----1
EXTENTS:
LOW-CCHH-----X'00040007'  LOW-RBA-----0          TRACKS-----1
HIGH-CCHH-----X'00040007'  HIGH-RBA-----10239
GDG BASE ----- MJK.GDG1
HISTORY
OWNER-IDENT-----OWNGDG      CREATION-----77.040      RCYVY-VOL-----333001      RCYVY-CI-----X'000015'
RELEASE-----2              EXPIRATION-----99.365      RCYVY-DEVT-----X'30502009'
ATTRIBUTES
LIMIT-----250          SCRATCH      EMPTY
ASSOCIATIONS
NONVSAN--MJK.GDG1.G0001V00

```

Figure 20 (Part 3 of 7). An Example of LISTCAT ALL Output

```

NONVSAM ---- MJK.GDG1.60001V00
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000016'
RELEASE-----2      EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001      DEVTYPE-----X'30502009'      FSEQN-----0
VOLSER-----333002      DEVTYPE-----X'30502009'      FSEQN-----0
ASSOCIATIONS
GDG-----MJK.GDG1
ALIAS----MJK.GDG1.ALIAS
ALIAS ----- MJK.GDG1.ALIAS
HISTORY
RELEASE-----2      RCVY-VOL-----333001      RCVY-DEVT----X'30502009'      RCVY-CI-----X'000017'
ASSOCIATIONS
NONVSAM--MJK.GDG1.60001V00
NONVSAM ----- MJK.NONVSAM1
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000018'
RELEASE-----2      EXPIRATION-----01.000      RCVY-DEVT----X'30502009'
VOLUMES
VOLSER-----333001      DEVTYPE-----X'30502009'      FSEQN-----0
ASSOCIATIONS----- (NULL)
CLUSTER ----- MJKCAT
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040
RELEASE-----2      EXPIRATION-----99.999
PROTECTION-PSWD----- (NULL)      RACF----- (NO)
ASSOCIATIONS
DATA-----VSAM.CATALOG.BASE.DATA.RECORD
INDEX----VSAM.CATALOG.BASE.INDEX.RECORD

```

Figure 20 (Part 4 of 7). An Example of LISTCAT ALL Output

```

DATA ----- VSAH.CATALOG.BASE.DATA.RECORD
HISTORY
OWNER-IDENT----- (NULL)          CREATION-----00.000
RELEASE-----2          EXPIRATION-----00.000
PROTECTION-PSWD----- (NULL)      RACF----- (NO)
ASSOCIATIONS
CLUSTER--MJKCAT
ATTRIBUTES
KEYLEN-----44          AVGLRECL-----505          BUFSPACE-----3072          CISIZE-----512
RKP-----0          MAXLRECL-----505          EXCPEXIT----- (NULL)      CI/CA-----40
SHROPTS(3,3) RECOVERY          SUBALLOC          NOERASE          INDEXED          NOWRITECHK          IMBED          NOREPLICAT
UNORDERED          NOREUSE          NONSPANNED          BIND          RECVABLE
STATISTICS
REC-TOTAL-----15          SPLITS-CI-----0          EXCPS-----24
REC-DELETED-----0          SPLITS-CA-----0          EXTENTS-----2
REC-INSERTED-----0          FREESPACE-%CI-----0          SYSTEM-TIMESTAMP:
REC-UPDATED-----0          FREESPACE-%CA-----0          X'0A51B6244C1B1000'
REC-RETRIEVED-----0          FREESPC-BYTES---381952
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----57          HI-ALLOC-RBA-----389120
SPACE-SEC-----18          HI-USED-RBA-----389120
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512          HI-ALLOC-RBA-----368640          EXTENT-NUMBER-4-----1
DEVTYPE-----X'30502009'          PHYRECS/TRK-----20          HI-USED-RBA-----20480          EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----3
LOW-KEY-----00
HIGH-KEY-----3F
HI-KEY-RBA-----6144
EXTENTS:
LOW-CCHH-----X'00000002'          LOW-RBA-----0          TRACKS-----54
HIGH-CCHH-----X'00020011'          HIGH-RBA-----368639
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512          HI-ALLOC-RBA-----389120          EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'          PHYRECS/TRK-----20          HI-USED-RBA-----389120          EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----3
LOW-KEY-----40
HIGH-KEY-----FF
HI-KEY-RBA-----368640
EXTENTS:
LOW-CCHH-----X'00030002'          LOW-RBA-----368640          TRACKS-----3
HIGH-CCHH-----X'00030004'          HIGH-RBA-----389119

```

Figure 20 (Part 5 of 7). An Example of LISTCAT ALL Output

INDEX ----- VSAM.CATALOG.BASE.INDEX.RECORD

```

HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2          EXPIRATION-----00.000
PROTECTION-PSWD----- (NULL)   RACF----- (NO)
ASSOCIATIONS
CLUSTER--MJKCAT
ATTRIBUTES
KEYLEN-----44          AVGLRECL-----0      BUFSIZE-----0      CISIZE-----512
RKP-----0            MAXLRECL-----505    EXCPEXIT----- (NULL)  CI/CA-----20
SHROPTNS(3,3) RECOVERY    SUBALLOC          NOERASE    NOWRITECHK      IMBED    NOREPLICAT      UNORDERED
NOREUSE          BIND
STATISTICS
REC-TOTAL-----3        SPLITS-CI-----0      EXCPS-----13       INDEX:
REC-DELETED-----0     SPLITS-CA-----0      EXTENTS-----3      LEVELS-----2
REC-INSERTED-----0    FREESPACE-%CI-----0  SYSTEM-TIMESTAMP:   ENTRIES/SECT-----7
REC-UPDATED-----0     FREESPACE-%CA-----0  X'8A51B6244C1B1000'  SEQ-SET-RBA-----30720
REC-RETRIEVED-----0   FREESPC-BYTES-----38912  HI-LEVEL-RBA-----0
ALLOCATION
SPACE-TYPE-----TRACK   HI-ALLOC-RBA-----40448
SPACE-PRI-----3       HI-USED-RBA-----40448
SPACE-SEC-----3
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512  HI-ALLOC-RBA-----30720  EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'  PHYRECS/TRK-----20  HI-USED-RBA-----512    EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME      TRACKS/CA-----1
EXTENTS:
LOW-CCHH-----X'00020012'  LOW-RBA-----0      TRACKS-----3
HIGH-CCHH-----X'00030001'  HIGH-RBA-----30719
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512  HI-ALLOC-RBA-----39936  EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'  PHYRECS/TRK-----20  HI-USED-RBA-----31232  EXTENT-TYPE-----X'80'
VOLFLAG-----PRIME      TRACKS/CA-----3
LOW-KEY-----00
HIGH-KEY-----3F
EXTENTS:
LOW-CCHH-----X'00000002'  LOW-RBA-----30720  TRACKS-----54
HIGH-CCHH-----X'00020011'  HIGH-RBA-----39935
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512  HI-ALLOC-RBA-----40448  EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'  PHYRECS/TRK-----20  HI-USED-RBA-----40448  EXTENT-TYPE-----X'80'
VOLFLAG-----PRIME      TRACKS/CA-----3
LOW-KEY-----40
HIGH-KEY-----FF
EXTENTS:
LOW-CCHH-----X'00030002'  LOW-RBA-----39936  TRACKS-----3
HIGH-CCHH-----X'00030004'  HIGH-RBA-----10447

```

Figure 20 (Part 6 of 7). An Example of LISTCAT ALL Output

```

VOLUME ----- 333001
HISTORY
RELEASE-----2          RCVY-VOL-----333001          RCVY-DEVT----X'30502009'          RCVY-CI-----X'000009'
CHARACTERISTICS
BYTES/TRK-----13165     DEVTYPE-----X'30502009'          MAX-PHYREC-SZ-----13030          DATASETS-ON-VOL-----5
TRKS/CYL-----19        VOLUME-TIMESTAMP:          MAX-EXT/ALLOC-----5          DATASPCS-ON-VOL-----1
CYLS/VOL-----411      X'8A51B6A76850E000'
DATASPACE
DATASETS-----5          FORMAT-1-DSCB:          ATTRIBUTES:
EXTENTS-----1          CCHHR-----X'0000000103'          SUBALLOC
SEC-ALLOC-----40        TIMESTAMP          EXPLICIT
TYPE-----TRACK          X'8A51B61D97B04000'          USERCAT
EXTENT-DESCRIPTOR:
TRACKS-TOTAL-----160     BEG-CCHH----X'00000002'          SPACE-MAP-----5947
TRACKS-USED-----89
DATASET-DIRECTORY:
DSN---MJKCAT          ATTRIBUTES----- (NULL)          EXTENTS-----3
DSN---MJK.CLUSTER.DATA  ATTRIBUTES----- (NULL)          EXTENTS-----1
DSN---MJK.CLUSTER1.INDEX  ATTRIBUTES----- (NULL)          EXTENTS-----1
DSN---MJK.ALT.INDEX1.DATA  ATTRIBUTES----- (NULL)          EXTENTS-----1
DSN---MJK.ALT.INDEX1.INDEX  ATTRIBUTES----- (NULL)          EXTENTS-----1

```

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

```

AIX -----1
ALIAS -----1
CLUSTER -----2
DATA -----3
GDG -----1
INDEX -----3
NONVSAM -----2
PAGESPACE -----0
PATH -----1
SPACE -----1
USERCATALOG -----0
TOTAL -----15

```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 20 (Part 7 of 7). An Example of LISTCAT ALL Output

LISTCAT ALLOCATION OUTPUT LISTING

When you specify the LISTCAT command and include the ALLOCATION parameter, each cataloged object with space allocated to it from a VSAM data space is listed (see Figure 21). All information about the object's space is listed, but none of the object's other cataloged information is listed. The entry types that can be specified when the ALLOCATION parameter is specified are limited to DATA and INDEX.

```

/* LIST SPACE ALLOCATION INFORMATION */

LISTCAT -
  ALLOCATION -
  CATALOG(MJKCAT)

AIX ----- MJK.ALT.INDEX1
HISTORY
  OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000010'
  RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

DATA ----- MJK.ALT.INDEX1.DATA
HISTORY
  OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000011'
  RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
ALLOCATION
  SPACE-TYPE-----TRACK
  SPACE-PRI-----6              HI-ALLOC-RBA-----73728
  SPACE-SEC-----6              HI-USED-RBA-----73728
VOLUME
  VOLSER-----333001            PHYREC-SIZE-----4096      HI-ALLOC-RBA-----73728      EXTENT-NUMBER-----1
  DEVTYPE-----X'30502009'      PHYRECS/TRK-----3          HI-USED-RBA-----73728      EXTENT-TYPE-----X'00'
  VOLFLAG-----PRIME            TRACKS/CA-----6
EXTENTS:
  LOW-CCHH----X'00040008'        LOW-RBA-----0              TRACKS-----6
  HIGH-CCHH---X'0004000D'        HIGH-RBA-----73727

INDEX ----- MJK.ALT.INDEX1.INDEX
HISTORY
  OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000012'
  RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
ALLOCATION
  SPACE-TYPE-----TRACK
  SPACE-PRI-----1              HI-ALLOC-RBA-----10240
  SPACE-SEC-----1              HI-USED-RBA-----512
VOLUME
  VOLSER-----333001            PHYREC-SIZE-----512        HI-ALLOC-RBA-----10240      EXTENT-NUMBER-----1
  DEVTYPE-----X'30502009'      PHYRECS/TRK-----20          HI-USED-RBA-----512        EXTENT-TYPE-----X'40'
  VOLFLAG-----PRIME            TRACKS/CA-----1
EXTENTS:
  LOW-CCHH----X'0004000E'        LOW-RBA-----0              TRACKS-----1
  HIGH-CCHH---X'0004000E'        HIGH-RBA-----10239

PATH ----- MJK.AIX1.PATH
HISTORY
  OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000014'
  RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

CLUSTER ----- MJK.CLUSTER1
HISTORY
  OWNER-IDENT-----OMNCLUST      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000E'
  RELEASE-----2              EXPIRATION-----77.040      RCVY-DEVT----X'30502009'

```

Figure 21 (Part 1 of 3). An Example of LISTCAT ALLOCATION Output

```

DATA ----- MJK.CLUSTER.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000000'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----2              HI-ALLOC-RBA-----24576
SPACE-SEC-----2              HI-USED-RBA-----24576
VOLUME
VOLSER-----333001          PHYREC-SIZE-----6144      HI-ALLOC-RBA-----24576      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'    PHYRECS/TRK-----2         HI-USED-RBA-----24576      EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----2
EXTENTS:
LOW-CCHH----X'00040005'      LOW-RBA-----0            TRACKS-----2
HIGH-CCHH---X'00040006'      HIGH-RBA-----24575
INDEX ----- MJK.CLUSTER1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000000F'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----1              HI-ALLOC-RBA-----10240
SPACE-SEC-----1              HI-USED-RBA-----512
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512       HI-ALLOC-RBA-----10240      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'    PHYRECS/TRK-----20        HI-USED-RBA-----512        EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----1
EXTENTS:
LOW-CCHH----X'00040007'      LOW-RBA-----0            TRACKS-----1
HIGH-CCHH---X'00040007'      HIGH-RBA-----10239
CLUSTER ----- MJKCAT
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040
RELEASE-----2              EXPIRATION-----99.999
DATA ----- VSAM.CATALOG.BASE.DATA.RECORD
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2              EXPIRATION-----00.000
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----57           HI-ALLOC-RBA-----389120
SPACE-SEC-----18           HI-USED-RBA-----389120
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512       HI-ALLOC-RBA-----368640      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'    PHYRECS/TRK-----20        HI-USED-RBA-----20480      EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----3
LOW-KEY-----00
HIGH-KEY-----3F
HI-KEY-RBA-----6144
EXTENTS:
LOW-CCHH----X'00000002'      LOW-RBA-----0            TRACKS-----54
HIGH-CCHH---X'00020011'      HIGH-RBA-----368639
VOLUME
VOLSER-----333001          PHYREC-SIZE-----512       HI-ALLOC-RBA-----389120      EXTENT-NUMBER-----1
DEVTYPE-----X'30502009'    PHYRECS/TRK-----20        HI-USED-RBA-----389120      EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME          TRACKS/CA-----3
LOW-KEY-----40
HIGH-KEY-----FF
HI-KEY-RBA-----368640
EXTENTS:
LOW-CCHH----X'00030002'      LOW-RBA-----368640        TRACKS-----3
HIGH-CCHH---X'00030004'      HIGH-RBA-----389119

```

Figure 21 (Part 2 of 3). An Example of LISTCAT ALLOCATION Output

INDEX ----- VSAM.CATALOG.BASE.INDEX.RECORD

```

HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2          EXPIRATION-----00.000
ALLOCATION
SPACE-TYPE-----TRACK
SPACE-PRI-----3        HI-ALLOC-RBA-----40448
SPACE-SEC-----3        HI-USED-RBA-----40448
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512   HI-ALLOC-RBA-----30720   EXTENT-NUMBER-----1
DEVTYPE-----X'30502009' PHYRECS/TRK-----20   HI-USED-RBA-----512    EXTENT-TYPE-----X'00'
VOLFLAG-----PRIME      TRACKS/CA-----1
EXTENTS:
LOW-CCHH-----X'00020012'  LOW-RBA-----0        TRACKS-----3
HIGH-CCHH-----X'00030001'  HIGH-RBA-----30719
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512   HI-ALLOC-RBA-----39936   EXTENT-NUMBER-----1
DEVTYPE-----X'30502009' PHYRECS/TRK-----20   HI-USED-RBA-----31232   EXTENT-TYPE-----X'80'
VOLFLAG-----PRIME      TRACKS/CA-----3
LOW-KEY-----00
HIGH-KEY-----3F
EXTENTS:
LOW-CCHH-----X'00000002'  LOW-RBA-----30720   TRACKS-----54
HIGH-CCHH-----X'00020011'  HIGH-RBA-----39935
VOLUME
VOLSER-----333001      PHYREC-SIZE-----512   HI-ALLOC-RBA-----40448   EXTENT-NUMBER-----1
DEVTYPE-----X'30502009' PHYRECS/TRK-----20   HI-USED-RBA-----40448   EXTENT-TYPE-----X'80'
VOLFLAG-----PRIME      TRACKS/CA-----3
LOW-KEY-----40
HIGH-KEY-----FF
EXTENTS:
LOW-CCHH-----X'00030002'  LOW-RBA-----39936   TRACKS-----3
HIGH-CCHH-----X'00030004'  HIGH-RBA-----40447

```

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

```

AIX -----1
ALIAS -----0
CLUSTER -----2
DATA -----3
GDG -----0
INDEX -----3
NONVSAM -----0
PAGESPACE -----0
PATH -----1
SPACE -----0
USERCATALOG -----0
TOTAL -----10

```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 21 (Part 3 of 3). An Example of LISTCAT ALLOCATION Output

LISTCAT HISTORY OUTPUT LISTING

When you specify the LISTCAT command and include the HISTORY parameter, only the name, ownerid, creation date, and expiration date are listed for each entry that is selected (see Figure 22). Only these types of entries have HISTORY information: ALTERNATEINDEX, CLUSTER, DATA, INDEX, NONVSAM, and PATH.

/* LIST HISTORY INFORMATION FOR EACH ENTRY */

LISTCAT -
HISTORY -
CATALOG(MJKCAT)

LISTING FROM CATALOG -- MJKCAT

```

AIX ----- MJK.ALT.INDEX1
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000010'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

DATA ----- MJK.ALT.INDEX1.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000011'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

INDEX ----- MJK.ALT.INDEX1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000012'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

PATH ----- MJK.AIX1.PATH
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000014'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

CLUSTER ----- MJK.CLUSTER1
HISTORY
OWNER-IDENT-----OWNCLUST      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000E'
RELEASE-----2              EXPIRATION-----77.140      RCVY-DEVT----X'30502009'

DATA ----- MJK.CLUSTER1.DATA
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000D'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

INDEX ----- MJK.CLUSTER1.INDEX
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'00000F'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

GDG BASE ----- MJK.GDG1
HISTORY
OWNER-IDENT-----OWNGDG      CREATION-----75.223      RCVY-VOL-----333001      RCVY-CI-----X'000015'
RELEASE-----2              EXPIRATION-----99.365      RCVY-DEVT----X'30502009'

NONVSAM ---- MJK.GDG1.G0001V00
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000016'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

ALIAS ----- MJK.GDG1.ALIAS
HISTORY
RELEASE-----2              RCVY-VOL-----333001      RCVY-DEVT----X'30502009'      RCVY-CI-----X'000017'

NONVSAM ----- MJK.NONVSAM1
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040      RCVY-VOL-----333001      RCVY-CI-----X'000018'
RELEASE-----2              EXPIRATION-----00.000      RCVY-DEVT----X'30502009'

```

Figure 22 (Part 1 of 2). An Example of LISTCAT HISTORY Output

```

CLUSTER ----- MJKCAT
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----77.040
RELEASE-----2          EXPIRATION-----99.999

DATA ----- VSAM.CATALOG.BASE.DATA.RECORD
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2          EXPIRATION-----00.000

INDEX ----- VSAM.CATALOG.BASE.INDEX.RECORD
HISTORY
OWNER-IDENT----- (NULL)      CREATION-----00.000
RELEASE-----2          EXPIRATION-----00.000

VOLUME ----- 333001
HISTORY
RELEASE-----2          RCVY-VOL-----333001    RCVY-DEVT----X'30502009'    RCVY-CI-----X'000009'

```

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

```

AIX -----1
ALIAS -----1
CLUSTER -----2
DATA -----3
GDG -----1
INDEX -----3
NONVSAM -----2
PAGESPACE -----0
PATH -----1
SPACE -----1
USERCATALOG -----0
TOTAL -----15

```

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 22 (Part 2 of 2). An Example of LISTCAT HISTORY Output

LISTCAT CREATION/EXPIRATION OUTPUT LISTING

When you specify the LISTCAT command and include the CREATION or EXPIRATION parameter (or both), entries that have a creation or expiration date are selected according to the number of days you specify in the subparameter.

In Figure 23, for example, because all entries in the listed catalog, MJKCAT, were created on the same day as the LISTCAT, no entries are listed as a result of the LISTCAT CREATION(5) job. When that job is run on an older catalog, each entry whose creation date is previous to the number of days specified with the CREATION parameter is listed (that is, the CREATION number of days specifies that all objects in the catalog at least 5 days old are to be listed). The creation date of the data and index objects of a cluster or alternate index is always the same as the creation date of its associated cluster or alternate index object. The creation date of the data and index object of a catalog is always set to 0.

When you list all entries of a catalog, and you specify the CREATION parameter, each user catalog connector entry and each alias entry are also listed regardless of their creation date.

When the LISTCAT EXPIRATION(20) job is run, each entry whose expiration date occurs within 20 days of today's date is listed. Because the expiration date of a cluster or alternate index is controlled by the cluster or alternate index object entry in the catalog, the expiration date for the data and index objects of a cluster or alternate index is always 0.

When you list all entries of a catalog and you specify the EXPIRATION parameter, each volume entry will be listed because volume entries have no expiration date.

These types of entries can have a creation or expiration date: ALTERNATEINDEX, PATH, CLUSTER, DATA, INDEX, NONVSAM, GDG, and PAGESPACE.

```
/* LIST EACH CATALOG ENTRY WHOSE CREATION DATE */
/* IS 5 DAYS AGO OR GREATER (THAT IS, THE OBJECT */
/* IS AT LEAST 5 DAYS OLD). */
```

```
LISTCAT -
  CREATION(5) -
  CATALOG(MJKCAT)
```

LISTING FROM CATALOG -- MJKCAT

```
ALIAS ----- MJK.GDG1.ALIAS
```

Figure 23 (Part 1 of 2). An Example of LISTCAT CREATION/EXPIRATION Output

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

AIX	-----	0
ALIAS	-----	1
CLUSTER	-----	0
DATA	-----	0
GDG	-----	0
INDEX	-----	0
NONVSAM	-----	0
PAGESPACE	-----	0
PATH	-----	0
SPACE	-----	0
USERCATALOG	-----	0
TOTAL	-----	1

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

/* LIST EACH CATALOG ENTRY WHOSE EXPIRATION DATE */
/* IS WITHIN 20 DAYS

LISTCAT -
EXPIRATION(20) -
CATALOG(MJKCAT)

LISTING FROM CATALOG -- MJKCAT

AIX ----- MJK.ALT.INDEX1
DATA ----- MJK.ALT.INDEX1.DATA
INDEX ----- MJK.ALT.INDEX1.INDEX
PATH ----- MJK.AIX1.PATH
ALIAS ----- MJK.GDG1.ALIAS
NONVSAM ----- MJK.NONVSAM1
VOLUME ----- 333001

LISTING FROM CATALOG -- MJKCAT

THE NUMBER OF ENTRIES PROCESSED WAS:

AIX	-----	1
ALIAS	-----	1
CLUSTER	-----	0
DATA	-----	1
GDG	-----	0
INDEX	-----	1
NONVSAM	-----	1
PAGESPACE	-----	0
PATH	-----	1
SPACE	-----	1
USERCATALOG	-----	0
TOTAL	-----	7

THE NUMBER OF PROTECTED ENTRIES SUPPRESSED WAS 0

IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0

IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0

Figure 23 (Part 2 of 2). An Example of LISTCAT CREATION/EXPIRATION Output

EXAMPLES OF LISTCAT IN A TSO ENVIRONMENT

The following examples illustrate the output produced at a TSO terminal for LISTCAT NAMES (default) and LISTCAT VOLUME. A TSO logon ID of IBMUSER is assumed.

For LISTCAT NAMES, the catalog name is printed followed by the names of all entries that have a high-level qualifier equal to the USER logon ID.

For LISTCAT VOLUME, all entrynames that have a high-level qualifier equal to the USER logon ID are printed followed by the volume serial numbers for those entries that contain volume information.

Note: Because volume serial numbers for a cluster or an alternate index are contained in the data and index components, the data and index must have been named on the initial DEFINE in order to list the volume serial numbers.

```
LOGON IBMUSER
```

```
READY  
LISTCAT
```

```
IN CATALOG: AMASTCAT  
IBMUSER.AIX  
IBMUSER.AIXDATA  
IBMUSER.AIXIDX  
IBMUSER.GDG  
IBMUSER.GDG.G0001V00  
IBMUSER.GDG.G0002V00  
IBMUSER.GDG.G0003V00  
IBMUSER.KSDS  
IBMUSER.KSDSDATA  
IBMUSER.KSDSIDX  
IBMUSER.NVSAM1  
IBMUSER.NVSAM2  
IBMUSER.NVSAM3  
IBMUSER.NVSAM4  
IBMUSER.NVSAM5  
READY
```

```
LISTCAT VOLUME
```

```
IBMUSER.AIX  
IBMUSER.AIXDATA  
--VOLUMES--  
333001  
IBMUSER.AIXIDX  
--VOLUMES--  
333001  
IBMUSER.GDG  
IBMUSER.GDG.G0001V00  
--VOLUMES--  
333001  
333002  
333003  
IBMUSER.GDG.G0002V00  
--VOLUMES--  
333004  
333005  
333006  
333007  
333008  
IBMUSER.GDG.G0003V00  
--VOLUMES--  
333009  
333010
```

IBMUSER.KSDS
IBMUSER.KSDSDATA
--VOLUMES--
333001
IBMUSER.KSDSIDX
--VOLUMES--
333001
IBMUSER.NVSAM1
--VOLUMES--
231401
231402
IBMUSER.NVSAM2
--VOLUMES--
231403
231404
231405
IBMUSER.NVSAM3
--VOLUMES--
231406
IBMUSER.NVSAM4
--VOLUMES--
231407
IBMUSER.NVSAM5
--VOLUMES--
231408
231409
231410
231411
231412
READY

APPENDIX C. JCL DD PARAMETERS TO TAKE CARE WITH

Because OS/VS does not disallow any DD statement parameters and subparameters for VSAM, you should be aware of the specifications that either do nothing in a VSAM environment or might cause problems for you. Figure 24 shows the DD statement parameters and subparameters to be avoided with VSAM.

Parameter	Subparameter	Comment
AFF	ddname	You must use this parameter carefully. If the cluster components, the data and its index, reside on unlike devices, the results of an AFF specification are unpredictable.
CHKPT	EOV	Because checkpoints at end of volume are executed only for BSAM and QSAM data sets, this parameter does not apply to VSAM data sets and need not be coded.
DATA		Because there is no way to get VSAM data into the input stream, this parameter is not applicable to VSAM.
DCB	All	The access-method control block (ACB), not the DCB, describes VSAM data sets; therefore, the DCB parameter is not applicable to VSAM. An access-method control block is generated by an ACB or GENCB macro, and can be modified by a MODCB macro.
DISP	CATLG	CATLGVSAM data sets are cataloged and uncataloged as a result of an Access Method Services command; if CATLG is coded, a message is issued, but the data set is not cataloged.
	DELETE	VSAM data sets are deleted as a result of an Access Method Services command; if DELETE is coded, a message is issued, but the data set is not deleted.
	MOD	For VSAM data sets, MOD is treated as if OLD were specified, except for processing with an ISAM program, in which case MOD indicates resume load.
	KEEP	Because KEEP is implied for VSAM data sets, it need not be coded.
	NEW	VSAM data spaces are initially allocated as a result of the Access Method Services DEFINE command. If NEW is specified, OS/VS also allocates space, and it is never used by VSAM. Moreover, an Access Method Services request for space may fail if the DISP=NEW acquisition of space causes too little space to remain available.

Figure 24 (Part 1 of 3). JCL DD Parameters

Parameter	Subparameter	Comment
	UNCATLG	VSAM data sets are cataloged and uncataloged as a result of Access Method Services commands; if UNCATLG is coded, a message is issued, but the data set is not uncataloged.
DSNAME	dsname(areaname)	The name is used; areaname is ignored.
	dsname(generation)	A generation data group entry must exist in the VSAM catalog for this to apply.
	dsname(member)	The name is used; member is ignored.
	All temporary dsnames	Because VSAM data sets are built by Access Method Services, which uses the data-set name supplied in the DEFINE command; temporary names cannot be used with VSAM.
	All backward DD references of the form *.ddname	If the object referred to is a cluster and the data set and index reside on unlike devices, the results of a backward DD reference are unpredictable.
LABEL	BLP, NL, NSL	Because these subparameters have no meaning for direct-access devices, they do not apply for VSAM data sets, which all reside on direct-access storage.
	IN	Because IN is used to override DCB subparameters and the DCB parameter does not apply to VSAM data sets, IN does not apply.
	OUT	Because OUT is used to override DCB subparameters and the DCB parameter does not apply to VSAM data sets, OUT does not apply.
LABEL	NOPWREAD	The password-protection bit is set for all VSAM data sets, regardless of the PASSWORD/NOPWREAD specification in the LABEL parameter.
	PASSWORD	The password-protection bit is set for all VSAM data sets, regardless of the PASSWORD/NOPWREAD specification in the LABEL parameter.
	SL, SUL	Although these parameters apply to direct-access storage devices, SL is always used for VSAM, whether you specify SL, SUL, or neither.
SEP	ddname	You must use this parameter carefully. If the cluster components, the data and its index, reside on unlike devices, the results of a SEP specification are unpredictable.
SPACE		VSAM data spaces are initially allocated as a result of the Access Method Services DEFINE command. If SPACE is specified, therefore, an extent is allocated that is never used by VSAM. Moreover, an Access Method Services request for space may fail as a result of the SPACE acquisition of space.

Figure 24 (Part 2 of 3). JCL DD Parameters

Parameter	Subparameter	Comment
SPLIT		VSAM data spaces are initially allocated as a result of the Access Method Services DEFINE command. If SPLIT is specified, OS/VS also allocates space, and it is never used by VSAM. Moreover, an Access Method Services request for space may fail if the SPLIT acquisition of space causes too little space to remain available.
SUBALLOC		VSAM data spaces are initially allocated as a result of the Access Method Services DEFINE command. If SUBALLOC is specified, OS/VS also allocates space, and it is never used by VSAM. Moreover, an Access Method Services request for space may fail if the SUBALLOC acquisition of space causes too little space to remain available. If the data set for which SUBALLOC is specified is a VSAM data set, the request is denied.
SYSOUT		If SYSOUT is coded with a mutually exclusive parameter (for example, DISP), the job step is terminated with an error message.
UCS	All	Because this parameter applies only to unit-record devices, it does not apply to VSAM.
UNIT	AFF	You must use this subparameter carefully. If the cluster components, the data and its index, reside on unlike devices, the results of UNIT=AFF are unpredictable.
	SEP	You must use this subparameter carefully. If the cluster components, the data and its index, reside on unlike devices, the results of UNIT=SEP are unpredictable.
VOLUME	REF	You must use this subparameter carefully. If the referenced volumes are not a subset of those contained in the catalog record for the data set, the results are unpredictable.
	volseq#	Results are unpredictable.
	volcount	This subparameter is used to request some number of nonspecific volumes. Because all VSAM volumes must be specifically defined before processing, volcount is not applicable to VSAM data sets.

Figure 24 (Part 3 of 3). JCL DD Parameters

APPENDIX D. INTERPRETING LISTCRA OUTPUT LISTINGS

When you code the LISTCRA command, you can specify options that allow you to tailor the contents of the LISTCRA output. This appendix illustrates the various types of the LISTCRA output, the order in which entries are listed, and the meanings of the listed fields.

There are five types of LISTCRA listings. Four of the types are illustrated in this appendix. The fifth type, SEQUENTIALDUMP, is exactly the same as the DUMP NOCOMPARE type except that the entries are not sorted into groups.

Each listed entry is identified by type (that is, cluster, nonVSAM, etc.) and by name. Entries are usually listed in alphabetic order within groups of entries, according to entryname (except for SEQUENTIALDUMP). However, if insufficient virtual storage is available for the sorting operation, the records are listed as they appear in the catalog recovery area.

On the listing, entries are sorted into three groups (except for a SEQUENTIALDUMP listing):

- VSAM entries (which are also sorted according to entryname). Cluster entries, alternate index entries, and their associated data, index, and path entries are within this group.
- Other entries that are sorted according to entryname. User catalog connector entries and nonVSAM entries are in this group.
- Unsorted entries: other entries that are not sorted according to entryname. Entries that are listed within this group depend on the type of LISTCRA listing requested.

The following list contains the abbreviation, type, and description for each kind of entry that can be listed as a result of the LISTCRA command:

Abbreviation	Type	Description
AIX	G	Alternate index entry
ALIA	X	Alias entry
CLUS	C	Cluster entry
DATA	D	Data component entry
FRSP	F	Freespace entry
GDGB	B	Generation data group entry
INDX	I	Index component entry
NONV	A	NonVSAM data set entry
OEXT	E	Extension record for an entry other than a volume entry
PATH	P	Path entry
UCAT	U	User catalog connector entry (in the master catalog only)
UPGD	Y	Upgrade set entry

Abbreviation	Type	Description
VEXT	W	Extension record for a volume entry
VOL	V	Volume entry

EXAMPLES OF LISTCRA LISTINGS

The five types of LISTCRA listings are:

- **NAME NOCOMPARE**—This is the default option. Each entryname, its volumes, and the name and volumes of each related entry are listed. Within the "unsorted entries" group, all catalog records not yet printed are printed. See Figure 25.
- **DUMP NOCOMPARE**—When DUMP is specified (NOCOMPARE is a default), each record is listed in the dump format. In addition, the name and volumes of each indirectly related entry are also listed. Within the "unsorted entries" group, the CRA's self-describing records (control intervals through 8) are printed, and any free space records and records not yet printed are printed. See Figure 26.
- **NAME COMPARE**—When COMPARE is specified (NAME is a default), only the name of each miscompared catalog entry is listed. A miscompared catalog entry is one whose information is not identical to the information contained in the entry's copy in the catalog recovery area. A MISCOMPARE message is printed that identifies the most serious level of miscomparison. Within the "unsorted entries" group, miscompared records not yet printed out are printed. See Figure 27.
- **DUMP COMPARE**—When DUMP and COMPARE are specified, only the records (in dump format) of each miscompared catalog entry are listed. A miscompared catalog entry is one whose information is not identical to the information contained in the entry's copy in the catalog area. For each entry, the catalog recovery area copy is listed first, followed by the catalog entry, followed by a line that contains asterisks to identify each miscompared byte. Following each entry, a MISCOMPARE message is printed that identifies the most serious level of miscomparison. Within the "unsorted entries" group, miscompared records not yet printed out are printed in the dump format. See Figure 28.

Note: As explained above, all MISCOMPARE messages result from a comparison between the catalog and the CRA if the comparison shows that the catalog and CRA records are not identical. No inference is intended as to which is correct. You must make this determination yourself by looking at other mismatches in the same listing and by examining related records in the catalog or CRA.

- **SEQUENTIALDUMP**—When SEQUENTIALDUMP is specified, each record in the catalog recovery area is printed. The format of the output is the same as the DUMP NOCOMPARE format, except that the entries are not sorted into groups or alphabetic sequence.

See OS/VS2 Catalog Management for a complete description of catalog recovery area record formats, catalog record formats, and relationships between catalog records.

JOB CONTROL LANGUAGE (JCL) FOR LISTCRA JOBS

The job control language (JCL) statements that can be used to list catalog recovery areas are:

```
//LISTCAT JOB ...
//STEP1 EXEC PGM=IDCAMS
//STPCAT DD DSN=YOURCAT,DISP=SHR
//CATDD DD DSN=YOURCAT,DISP=OLD
//CRADD1 DD VOL=SER=333001,UNIT=3330,DISP=OLD
//CRADD2 DD VOL=SER=333002,UNIT=3330,DISP=OLD
//OUTDD DD DSN=LISTCRA.OUTPUT,UNIT=2400-3,
// VOL=SER=TAPE10,LABEL=(1,NL),DISP=(NEW,KEEP),
// DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629)
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
LISTCRA -
COMPARE -
DUMP -
INFILE(CRADD1,CRADD2) -
CATALOG(YOURCAT/PASSWORD CATDD) -
MASTERPW(SECRET) -
OUTFILE(OUTDD)
/*
```

The JOB statement contains user and accounting information required for your installation.

The EXEC statement identifies the program to be executed, IDCAMS (that is, the Access Method Services program).

- STEPCAT DD, which allocates your catalog. This is only required if the COMPARE option is specified. The catalog must be open as a catalog before it can be opened as a data set through the DD card that so identifies it for the comparisons.
- CATDD, which specifies the catalog to be opened as a data set and compared to the catalog recovery areas. This is only required if the COMPARE option is specified.
- CRADD1, which specifies a volume whose catalog recovery area (CRA) is to be listed.
- CRADD2, which specifies another volume whose CRA is to be listed.
- OUTDD, which specifies an alternate output file, so that the LISTCRA output can be written onto an auxiliary storage device. The LISTCRA command's OUTFILE parameter points to the OUTDD DD statement. Only the LISTCRA output is written to the alternate output device. JCL statements, system messages, and job statistics are written to the SYSPRINT output device.
 - DSN=LISTCRA.OUTPUT specifies the name for the magnetic tape file.
 - UNIT=2400-3 and VOL=SER=TAPE10 specifies that the file is to be contained on magnetic tape volume TAPE10.
 - LABEL=(1,NL) specifies that this is the first file on a nonlabeled tape. You can also use a standard-labeled tape by specifying LABEL=(1,SL). If subsequent job steps produce additional files of LISTCRA output on the same tape volume, you should increment the file number in each job step's LABEL subparameter (that is, LABEL=(2,NL) for the second job step, LABEL=(3,NL) for the third job step, etc.).
 - DISP=(NEW,KEEP) specifies that this is a new tape file and is to be rewound when the job finishes. If a subsequent job step prints the tape, DISP=(NEW,PASS) should be specified. If your job step contains more than

one LISTCRA command, DISP=(MOD,KEEP) or DISP=(MOD,PASS) can be used to concatenate all of the LISTCRA output in one sequential file.

- DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629) specifies that the LISTCRA output records are variable-length, blocked 5-to-1, and are preceded by an ANSI print-control character.

- SYSPRINT DD, which is required for each Access Method Services job step. It identifies the output queue, SYSOUT=A, on which all LISTCRA output and system output messages are printed (unless the OUTFILE parameter and its associated DD statement is specified—see OUTDD above).

Note: If you want all output to be written to an auxiliary storage device, replace 'OUTDD' with 'SYSPRINT' in the OUTDD DD statement and omit the SYSPRINT DD SYSOUT=A statement.

- SYSIN DD, which specifies, with an asterisk (*), that the statements that follow are the input data statements. A '/*' terminates the input data statements.

The LISTCRA command parameters shown above are common to the LISTCRA DUMP COMPARE example that follows. Other LISTCRA parameters may be coded and the output that results is illustrated.

COMPARE, which specifies that the CRA entries are to be compared with the catalog entries identified by the CATALOG parameter. Only those that miscompare will be listed.

DUMP, which specifies that each listed entry is to be printed in its entirety in both hexadecimal and character form.

INFILE, which specifies the two DD statements, CRADD1 and CRADD2, which identify the two volumes whose CRA's are to be compared with the catalog entries and the miscomparing entries listed.

CATALOG, which identifies the DD statement, CATDD, which identifies the catalog, YOURCAT, whose entries are to be compared with those in the CRAs. If the catalog is password protected, its master password, PASSWORD, is required.

MASTERPW, which specifies the master password of the master catalog, SECRET, to enable the CRAs to be OPENED.

OUTFILE, which points to the OUTDD DD statement. The OUTDD DD statement allocates an alternate output file for the LISTCRA output.

If you want to print the LISTCRA output that is contained on an alternate output file, you can use the IEBGENER program. The following shows the JCL required to print the alternate output file, LISTCRA.OUTPUT, that was allocated previously:

```
//PRINTOUT JOB ...
//STEP1 EXEC PGM=IEBGENER
//SYSUT1 DD DSN=LISTCRA.OUTPUT,UNIT=2400-3,
          DCB=(RECFM=VBA,LRECL=125,BLKSIZE=629)
//SYSUT2 DD SYSOUT=A
//SYSPRINT DD SYSOUT=A
//SYSIN DD DUMMY
/*
//
```

LISTCRA OUTPUT LISTING

When you specify LISTCRA without any parameters, each entryname, its volume(s), and the name and volume(s) of each related entry are listed. The same listing would result if the NAME and NOCOMPARE parameters were specified.

You can use this type of listing to list the name of each catalog entry whose copy is in the catalog recovery area, and to determine the number of entries in each catalog recovery area. The total number of entries is an approximate size, in records, of the catalog recovery area.

LISTING OF CATALOG RECOVERY AREA FOR VOLUME -- 333301 -- VSAM ENTRIES

VOL - 333301
CRAVOLRCD - 10/07/74 23:35:00
F4DSCBVSAM - 10/07/74 23:35:00
F4DSCBDUMP - 10/07/74 23:35:00

CLUS - AA.LISTCRA.ESDS

DATA - TF41C9A0.VSAMDSSET.DFD74280.T861DAD4.TF41C9A0
DATA VOL -
333301

AIX - AA2.LR.ESDS
DATA VOL -
333301
INDX VOL -
333301

AIX - AA1.LR.ESDS
DATA VOL -
333301
INDX VOL -
333301

PATH - AAU.LR.ESDS

UPGD -

CLUS - AA.LISTCRA.KSDS

DATA - T5C55DD0.VSAMDSSET.DFD74280.T861DB0A.T5C55DD0
DATA VOL -
333301

INDX - T5C56E70.VSAMDSSET.DFD74280.T861DB0A.T5C56E70
INDX VOL -
333301

AIX - AA1.LR.ESDS

DATA - TCD41020.VSAMDSSET.DFD74280.T861DB23.TCD41020
DATA VOL -
333301

INDX - TCD41F00.VSAMDSSET.DFD74280.T861DB23.TCD41F00
INDX VOL -
333301

CLUS - AA.LISTCRA.ESDS

PATH - AA1U.LR.ESDS

PATH - AA1N.LR.ESDS

CLUS - LR.MCKEYRNG.KSDS

DATA - TF081480.VSAMDSSET.DFD74280.T861DB0E.TF081480
DATA VOL - HIGH KEY
333301 - C1C1C1C1C1C1C1C1C1C5F9
333301 - C1C1C1C1C1C1C1C1C1E9E9

INDX - TF082370.VSAMDSSET.DFD74280.T861DB0E.TF082370
INDX VOL -
• 333301
333301

Figure 25 (Part 1 of 2). An Example of LISTCRA NAME NOCOMPARE Output

LISTING OF CATALOG RECOVERY AREA FOR VOLUME -- 333301 -- OTHER ENTRIES

UCAT - USERCAT2
DATA VOL -
333304

NUMBER OF ENTRIES PROCESSED

CLUS - 3
DATA - 4
AIX - 1
INDX - 2
PATH - 3
VOL - 1
UPGD - 1
SUM - 15

LISTCRA NAME NOCOMPARE Output Notes:

1. Volume serial number
2. Identifies the group that the following entries are within (in this case, VSAM Entries).
3. The time stamps:

CRAVOLRCD is the timestamp of the CRA volume record. It is updated only when the first catalog record and first CRA volume record are updated, after the CRA is opened when VSAM is running under OS/VS. (Note: If the volume is moved to a DOS/VS system and used there, the CRAVOLRCD timestamp is updated each time the catalog and CRA volume records are updated for a space allocation change.)

F4DSCBVSAM is a timestamp in the format-4 DSCB, and is updated in the same manner as the CRAVOLRCD timestamp.

F4DSCBDUMP is a timestamp in the format-4 DSCB, and is updated in the same manner as the CRAVOLRCD timestamp. It is also updated whenever the OS/VS utility program IEHDASDR dumps the volume.
4. VSAM entries are listed along with the entryname and volume(s) of each related entry.
5. Paths are shown as related only to the entry that the path provides access to. A path that serves as an alias for a VSAM data set is listed with the data-set's entry. A path that shows the relationship of an alternate index to a base cluster is listed with the alternate-index's entry only.
6. UPGD indicates that there are alternate indexes in the base-cluster's upgrade set. If you want to identify the indexes in the upgrade set, use the DUMP NOCOMPARE option.
7. The high-key value of each key range is shown.
8. SUM is the total number of catalog recovery area entries printed.

Figure 25 (Part 2 of 2). An Example of LISTCRA NAME NOCOMPARE Output

LISTCRA DUMP OUTPUT LISTING

When you specify LISTCRA with the DUMP parameter, each record in the catalog recovery area is listed in dump format. In addition, the name and volumes of each indirectly related entry are also listed. The same listing would result if the DUMP and NOCOMPARE parameters were specified.

When you specify LISTCRA SEQUENTIALDUMP, the same listing results—except that the records are listed in their entry sequence within the catalog recovery area.

You can use this type of listing to list the complete contents of each catalog entry whose copy is in the catalog recovery area, and to determine the exact number of entries and records in each catalog recovery area.

LISTING OF CATALOG RECOVERY AREA FOR VOLUME -- 333301 -- VSAM ENTRIES

```

1 VOL - 333301 2
000009 0000000C 01F3F3F3 F3F0F100 00093050 2009861D 8A4E0000 00000000 00000000 .....333301.....6.....+.....
0020 00000000 00000000 00000000 00000000 E301BD00 78F3F3F3 F3F0F100 00000000 00000000 .....L.V..... 333301.....
0040 00000000 00000000 00000000 00000000 00000000 00000000 00000000 00861DEE .....
0060 1152A820 00305020 09000032 E6019B00 13336DBF BF380002 00000500 00000000 .....6.....W.....
0080 00000000 2600003E 00000000 000D8500 0100000E 85000200 000F8500 03000037 .....
00A0 86000100 00378800 01000037 88000200 00378800 03000037 88000400 00378800 .....
00C0 05000037 88000600 00480800 07000037 88000800 00378800 09000037 88000A00 .....
00E0 00378800 0B000037 88000C00 00378800 0D000037 88000E00 00378800 0F000037 .....
0100 88001000 00378800 11000037 88001200 00378800 1300003E 88001400 003E8800 .....
0120 1500003E 88001600 003E8800 1700003E 88001800 003E8800 1900003E 88001A00 .....
0140 003E8800 1B000000 08001C00 000C0800 1D000018 08001E00 00240800 1F000030 .....
0160 08002000 003C0800 21000C00 0051861D B1D60000 002861D B1D60000 .....O.....O.....
0180 46000C00 0054861D B2050000 47000C00 0055861D B2050000 48000C00 0057861D .....
01A0 B23F0000 4C000C00 0058861D B23F0000 50000C00 0060861D B79C0000 2C1DB173 .....).f.....f.....

```

```

6 CLUS - KSDS01
00000E 00000017 01F3F3F3 F3F0F100 000E3050 2009861D 8A4E0000 00000000 00000000 .....333301.....6.....+.....
0020 00000000 00000000 00000000 C300A800 60D2E2C4 E2F0F140 40404040 40404040 .....C.....(KSDS01.....
0040 40404040 40404040 40404040 40404040 40404040 40404040 40404040 40FF0000 .....
0060 FFFFFFFF FF74280F 00000F00 00000000 00070000 00C00000 00000201 00006024 .....
0080 0200000C 02030000 12020400 00004401 0006C400 00180006 C9000019 0006D900 .....D.....I.....R.....
00A0 001A0006 C700001B 00000000 00000000 00000000 00000000 00000000 00000000 .....G.....
DATA - T188D0C0.VSAMDSSET.DFD74280.T861D99E.T188D0C0
00000D 00000018 01F3F3F3 F3F0F100 000D3050 2009861D 8A4E861D 99E20000 00000000 .....333301.....6.....+.....S.....
0020 00000000 00000000 00000000 C4016200 8FE3F1F8 F8C4F0C3 F04BE5E2 C1D4C4E2 .....D.....T188D0C0.VSAMDS
0040 C5E34BC4 C6C4F7F4 F2F8F04B E3F8F6F1 C4F9F9C5 4BE3F1F8 F8C4F0C3 F0FFFFFF ET.DFD74280.T861D99E.T188D0C0...
0060 FFFFFFFF FF74280F 00000000 20000000 22000000 01000001 80000000 00000030 .....
0080 0000000F F90000FF FFFFFFFF FFFFFFF0 00000000 06000000 C0000000 00010100 .....9.....
00A0 00620201 0000A902 02000068 03010000 00440100 62608000 60000000 03001400 .....
00C0 00000300 00000000 00000010 0000000F F9000000 00000000 00000000 00000000 .....9.....
00E0 00000000 00000000 00000000 00000000 01000000 00000000 00000000 00000000 .....
0100 00000000 00000030 00000000 00000000 00000000 000006C3 00001703 27305020 .....C.....6.....
0120 09F3F3F3 F3F0F100 00800100 00000000 00000000 00300000 00100000 03000140 .....333301.....
0140 00010000 00000014 00010002 00000002 00000001 00000000 00002FFF 0006E800 .....Y.....
0160 00150000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 .....
INDX - T188E190.VSAMDSSET.DFD74280.T861D99E.T188E190
00000F 00000019 01F3F3F3 F3F0F100 000F3050 2009861D 8A4E861D 99E20000 00000000 .....333301.....6.....+.....S.....
0020 00000000 00000000 00000000 C9015700 8FE3F1F8 F8C5F1F9 F04BE5E2 C1D4C4E2 .....I.....T188E190.VSAMDS
0040 C5E34BC4 C6C4F7F4 F2F8F04B E3F8F6F1 C4F9F9C5 4BE3F1F8 F8C5F1F9 F0FFFFFF ET.DFD74280.T861D99E.T188E190...
0060 FFFFFFFF FF74280F 00000000 2000FFFF FFFF0000 01000001 80000000 00000028 .....
0080 00FFFFFF FF0000FF FFFFFFFF FFFFFFF0 00000000 05000000 C0000000 00010100 .....
00A0 00620201 00006803 01000000 44010062 60000060 00010003 00140000 00140000 .....
00C0 00000000 00000200 000001F9 00000000 00000000 00000000 00000000 00000000 .....9.....
00E0 00000000 00000000 00000001 00000000 00000000 00000000 00000000 00000000 .....
0100 00002800 00000000 00000000 00000000 0006C300 00170327 30502009 F3F3F3F3 .....C.....6.....3333
0120 F0F10000 80010000 00000000 00000000 28000000 02000014 00140000 02000000 01.....
0140 00001400 01000200 01000200 01000100 00000000 0027FF00 00000000 00000000 .....
PATH - PATH01
000010 0000001A 01F3F3F3 F3F0F100 00100000 0000861D 8A4E0000 00000000 00000000 .....333301.....6.....+.....
0020 00000000 00000000 00000000 D9009800 6CD7C1E3 C8F0F140 40404040 40404040 .....R.....(PATH01.....
0040 40404040 40404040 40404040 40404040 40404040 40404040 40FF0000 .....
0060 FFFFFFFF FF74280F 00000F80 00000000 00000000 00C00000 00000201 00000602 .....
0080 0200000C 02030006 C3000017 0006C400 00180006 C9000019 00000000 00000000 .....C.....D.....I.....

```

Figure 26 (Part 1 of 2). An Example of LISTCRA DUMP NOCOMPARE Output

```

AIX - AIX01
DATA VOL -
  333301
INDX VOL -
  333301
UPGD -
000014 00000015 01F3F3F3 F3F0F100 00143050 2009861D 8A4E0000 00000000 00000000 .....333301.....6.....+.....
0020 00000000 00000000 00000000 E8004B00 31000000 00000200 0000C000 00000002 .....
0040 01000000 0A00001C03 00001B00 00000000 00000000 00000000 00000000 .....
DATA - TF081480.VSAMDSSET.DFD74280.T861DBOE.TF081480
000038 00000048 01F3F3F3 F3F0F100 00383050 2009861D 8A4E861D B0F00000 00000000 .....333301.....6.....+...0.....
0020 00000000 00000000 00000000 C401C500 8FE3C6F0 F8F1F4F8 F04BE5E2 C1D4C4E2 .....D.E..TF081480.VSAMDS
0040 C5E34BC4 C6CAF7F4 F2F8F04B E3F8F6F1 C4C2F0C5 4BE3C6F0 F8F1F4F8 F0FFFFF .....ET.DFD74280.T861DBOE.TF081480...
0060 FFFFFFFF FF74280F 00000000 20000000 06000000 01000001 80000000 00000050 .....6
0080 00000000 500000FF FFFFFFFF FFFFFFF0 00000000 06000000 C0000000 00010100 .....6
00A0 00620201 00006803 010000BD 03020000 00440100 62608400 60000000 0A000A00 .....
00C0 00001400 00000000 00000002 00000000 50000000 00000000 00000000 00000000 .....6
00E0 00000000 00000000 00000000 00000000 02000000 00000000 00000000 00000000 .....6
0100 00000000 00000050 00000000 00000000 00000000 000006C3 00004703 27305020 .....6
0120 09F3F3F3 F3F0F100 00800100 00000000 00000000 00280000 00020000 14000140 .....333301.....6
0140 0018000A C1C1C1C1 C1C1C1C1 C1C1000A C1C1C1C1 C1C1C1C1 C5F90014 00010003 .....AAAAAAAAAAAAAAAAA9.....
0160 00040003 00040001 00000000 000027FF 03273050 2009F3F3 F3F3F0F1 00008001 .....6.333302
0180 00000000 00002800 00005000 00000200 00140009 40001800 0AC1C1C1 C1C1C1C1 .....6
01A0 C1C6C100 0AC1C1C1 C1C1C1C1 C1E9E900 14000100 03000500 03000500 01000028 .....APA.AAAAAAAAAAZZ.....
01C0 0000004F FF000000 00000000 00000000 00000000 00000000 00000000 00000000 .....
INDX - TF082370.VSAMDSSET.DFD74280.T861DBOE.TF082370
00003A 00000049 01F3F3F3 F3F0F100 003A3050 2009861D 8A4E861D B0F00000 00000000 .....333301.....6.....+...0.....
0020 00000000 00000000 00000000 C9018900 8FE3C6F0 F8F2F3F7 F04BE5E2 C1D4C4E2 .....I..TF082370.VSAMDS
0040 C5E34BC4 C6CAF7F4 F2F8F04B E3F8F6F1 C4C2F0C5 4BE3C6F0 F8F2F3F7 F0FFFFF .....ET.DFD74280.T861DBOE.TF082370...
0060 FFFFFFFF FF74280F 00000000 2000FFFF FFFF0000 01000001 80000000 00000028 .....
0080 00FFFFF FF0000FF FFFFFFFF FFFFFFF0 00000000 06000000 C0000000 00010100 .....
00A0 00620201 00006803 010000A9 03020000 00440100 62600400 60000400 0A000A00 .....
00C0 00001400 00000000 00000002 00000001 F9000000 00000000 00000000 00000000 .....9
00E0 00000000 00000000 00000000 00000000 01000000 00000000 00000000 00000000 .....0.....6
0100 00000000 00000028 00000000 00000000 00000000 000006C3 00004703 27305020 .....
0120 09F3F3F3 F3F0F100 00800100 00000000 00000000 00280000 00020000 14000140 .....333301.....6
0140 00190000 00000014 00010003 00060003 00060001 00000000 000027FF 03273050 .....6
0160 2009F3F3 F3F3F0F1 00004000 00000000 00000000 00000000 00000200 00140001 .....333301.....6
0180 40001900 00000000 00000000 00000000 00000000 00000000 00000000 00000000 .....

```

LISTCRA DUMP NOCOMPARE Output Notes:

Note: When you specify SEQUENTIALDUMP, the same information is listed. The entries are not sorted according to type of entry, or according to alphameric-sequence of the entryname field. See also Note 7 below.

1. The 3-byte CRA control interval number, in hexadecimal form.
2. The 3-byte control interval number of the catalog record for which this is a copy, in hexadecimal form.
3. The entrytype at offset X'2C' (decimal 44), a character that identifies the type of catalog entry being listed.
4. A 44-byte field at offset X'31' (decimal 49) that contains the entry's entryname (padded with binary zeros), or an 8-byte field followed by 36 bytes of binary zero that contains the volume entry's volume serial number.
5. The volume serial number of the recovery volume.
6. A 2-byte displacement value. Printing is suppressed after the last line containing data is listed, even though all catalog records are 512 bytes long.
7. As in the NAME NOCOMPARE output, the alternate index name and the fact that there is an upgrade set is given, along with the volume serial numbers of the alternate index's data and index components. (**Note:** If SEQUENTIALDUMP is specified, the information described with this note is not listed.)
8. The first byte identifies the alternate index's data component and the next three bytes contain its catalog control interval number.
9. The first byte identifies the alternate index's index component and the next three bytes contain its catalog control interval number.
10. Since there are two key ranges for this cluster, there are two volume information sets of fields in the cluster's data and index component entries.

Figure 26 (Part 2 of 2). An Example of LISTCRA DUMP NOCOMPARE Output

LISTCRA COMPARE OUTPUT LISTING

When you specify LISTCRA with the COMPARE parameter, each record in the catalog recovery area is compared to its original in the catalog. The name of each catalog entry that miscompares is listed. A miscompared catalog entry is one whose information is not identical to the information contained in the entry's copy in the catalog recovery area. The same listing would result if the NAME and COMPARE parameters were specified.

You can use this type of listing to determine the number of damaged entries in your catalog, and you can get an idea of the type of damage that occurred to each entry. You can issue the EXPORTRA and IMPORTRA commands to recover the catalog and make its damaged entries usable.

LISTING OF CATALOG RECOVERY AREA FOR VOLUME -- 333301 -- VSAM ENTRIES

CATVOLRCD - 09/17/42 23:53:47
CRAVOLRCD - 10/07/74 23:35:00
F4DSCBVSAM - 10/07/74 23:35:00
F4DSCBDUMP - 10/07/74 23:35:00

CLUS - LR.DELETED.ESDS

* MISCOMPARES - CATALOG ENTRY HAS DIFFERENT NAME

CLUS - LR.MCHURBA.ESDS

DATA - T73FEDA0.VSAMDSSET.DFD74280.T861DAE1.T73FEDA0
DATA VOL -
333301

* MISCOMPARES - HIGH USED RBA

CLUS - LR.MCKEYRNG.KSDS

DATA - TF081480.VSAMDSSET.DFD74280.T861DB0E.TF081480
DATA VOL - HIGH KEY
333301 - C1C1C1C1C1C1C1C5F9
333301 - C1C1C1C1C1C1C1E9E9

* MISCOMPARES - HIGH USED RBA

INDX - TF082370.VSAMDSSET.DFD74280.T861DB0E.TF082370
INDX VOL -
333301
333301

* MISCOMPARES - STATISTICS

NUMBER OF ENTRIES PROCESSED

CLUS - 23
DATA - 18
AIX - 8
INDX - 11
PATH - 5
VOL - 1
UPGD - 3
SUM - 69

IDC0665I NUMBER OF ENTRIES THAT MISCOMPARED IN THIS CRA - 3

IDC0877I NUMBER OF RECORDS THAT MISCOMPARED IN THIS CRA - 4

LISTCRA NAME COMPARE Output Notes:

1. In addition to the timestamps described for the NAME NOCOMPARE output (noted previously), the timestamp in the catalog volume record, CATVOLRCD, is listed. CATVOLRCD is updated in the same manner as the CRAVOLRCD timestamp.
2. The MISCOMPARES message always refers to the record listed above it. See "Regaining Access to Data" under "Catalog Recovery" in the chapter on Data Security and Protection for the cause and seriousness of this MISCOMPARES message.
3. The number of records is sometimes larger than the number of entries, because an entry might consist of a catalog record and one or more extension records that failed to compare correctly.

Figure 27. An Example of LISTCRA NAME COMPARE Output

LISTCRA DUMP COMPARE OUTPUT LISTING

When you specify LISTCRA with the DUMP and COMPARE parameters, each record in the catalog recovery area is compared to its original in the catalog. Each catalog entry's copy in the catalog recovery area that mismatches is listed, followed by its original catalog entry (which is damaged), followed by asterisks to indicate each miscompared byte.

You can use this type of listing to determine the exact damage that occurred to each entry. You might be able to correct some or all of the damage by using the ALTER command. You can issue the EXPORTRA and IMPORTRA commands to recover the catalog and make its damaged entries usable.

```

LISTING OF CATALOG RECOVERY AREA FOR VOLUME -- 333301 -- VSAM ENTRIES

CLUS - LR.MCHURBA.ESDS
1 -> 000020 0000002C 01F3F3F3 F3F0F100 00203050 2009861D 8A4E0000 00000000 00000000 .....333301.....6.....+.....
CATRCD 0000002C 01F3F3F3 F3F0F100 00203050 2009861D 8A4E0000 00000000 00000000 .....333301.....6.....+..... 2
    0020 00000000 00000000 00000000 C3008200 6CD3D94B D4C3C8E4 D9C2C14B C5E2C4E2 .....C....(LR.MCHURBA.ESDS
    00000000 00000000 00000000 C3008200 6CD3D94B D4C3C8E4 D9C2C14B C5E2C4E2 .....C....(LR.MCHURBA.ESDS
    0040 40404040 40404040 40404040 40404040 40404040 40404040 40404040 40FFFFFFF .....
    40404040 40404040 40404040 40404040 40404040 40404040 40404040 40FFFFFFF .....
    0060 FFFFFFFF FF74280F 00000F00 00000000 00000000 00020100 00004401 0006C400 .....D.
    FFFFFFFF FF74280F 00000F00 00000000 00000000 00020100 00004401 0006C400 .....D.
    0080 002D0000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 .....
    002D0000 00000000 00000000 00000000 00000000 00000000 00000000 00000000 .....

DATA - T73FEDA0.VSAMDSSET.DFD74280.T861DAE1.T73FEDA0
000021 0000002D 01F3F3F3 F3F0F100 00213050 2009861D 8A4E861D AE2E0000 00000000 .....333301.....6.....+.....
CATRCD 0000002D 01F3F3F3 F3F0F100 00213050 2009861D 8A4E861D AE2E0000 00000000 .....333301.....6.....+.....
    0020 00000000 00000000 00000000 C4015700 8FE4F7F3 C6C5C4C1 F04BE5E2 C1D4C4E2 .....D....T73FEDA0.VSAMDS
    00000000 00000000 00000000 C4015700 8FE4F7F3 C6C5C4C1 F04BE5E2 C1D4C4E2 .....D....T73FEDA0.VSAMDS
    0040 C5E34BC4 C6C4F7F4 F2F8F04B E3F8F6F1 C4C1C5F1 4BE3F7F3 C6C5C4C1 F0FFFFFFF ET.DFD74280.T861DAE1.T73FEDA0...
    C5E34BC4 C6C4F7F4 F2F8F04B E3F8F6F1 C4C1C5F1 4BE3F7F3 C6C5C4C1 F0FFFFFFF ET.DFD74280.T861DAE1.T73FEDA0...
    0060 FFFFFFFF FF74280F 00000000 20800000 04000000 01000001 80000000 00000028 .....
    FFFFFFFF FF74280F 00000000 20800000 04000000 01000001 80000000 00000028 .....
*
* * * * *
0080 00000000 500000FF FFFFFFFF FFFFFFFF00 00000000 05000000 00000000 00010100 .....6..... 3
00000000 500000FF FFFFFFFF FFFFFFFF00 00000000 05000000 00000000 00010100 .....6.....
00A0 00620201 00006803 01000000 44010062 60000060 00000000 00000000 00140000 .....
00620201 00006803 01000000 44010062 60000060 00000000 00000000 00140000 .....
* MISCOMPARES - HIGH USED RBA 4

```

LISTCRA DUMP COMPARE Output Notes:

1. The 3-byte CRA control interval number in hexadecimal form.
2. Each pair of hexadecimal-data lines are the information in the catalog recovery area's copy of the catalog record, followed by the information in the catalog record, for the same displacement.
3. Asterisks are printed below the bytes in which a miscompare exists and, in the left margin, the miscompared lines are flagged with a single asterisk.
4. See "Regaining Access to Data" under "Catalog Recovery" in the chapter on Data Security and Protection for the cause and seriousness of this MISCOMPARES message.

Figure 28. An Example of LISTCRA DUMP COMPARE Output

APPENDIX E. SAMPLE OUTPUT FROM CHKLIST

Figure 29 shows the format of output from the CHKLIST command.

For each checkpoint entry listed under "CHECKID," the other columns give dsname, ddname, unit, and volume information for each tape data set that was open at the time of the checkpoint. Asterisks under "UNIT" indicate an unrecognizable unit type. "VOLUME X OF Y IS CURRENT" indicates the volume sequence number of the volume mounted at the time of the checkpoint and the total number of volumes in the data set. Each volume's volume serial number is listed, and an asterisk indicates the volume mounted at the time of the checkpoint.

```
CHKLIST INFILE(CHKPT) CHECKID(C0000001 C0000002)
IDCAMS SYSTEM SERVICES          TIME: 21:34:43      01/16/75
OPEN TAPE DATA SET LIST FROM CHECKPOINT DATA SET - CHKPT DATASET
CHECKID      DSNAME              DDNAME  UNIT          VOLUMES - * INDICATES CURRENT VOLUME
C0000001
      USER.TAPE.DATASET0  DDN2VS11 2400-7TRK VOLUME  1 OF  1 IS CURRENT
                        120001*
      USER.TAPE.DATASET1  DDN3VS11 2400-9TRK VOLUME  3 OF  4 IS CURRENT
                        130001 130002 130003* 130004
      USER.TAPE.DATASET2  DDN4VS11 3400-7TRK VOLUME  1 OF  1 IS CURRENT
                        140001*
C0000002
      USER.TAPE.DATASET3  DDNAME32 2400-9TRK VOLUME  8 OF  8 IS CURRENT
                        230001 230002 230003 230004 230005
                        230006 230007 230008*
      USER.TAPE.DATASET4  DDNAME42 2400-9TRK VOLUME  20 OF 24 IS CURRENT
                        240001 240002 240003 240004 240005
                        240006 240007 240008 240009 240010
                        240011 240012 240013 240014 240015
                        240016 240017 240018 240019 240020*
                        240021 240022 240023 240024
IDC0001I FUNCTION COMPLETED, HIGHEST CONDITION CODE WAS 0
IDC0002I IDCAMS PROCESSING COMPLETE. MAXIMUM CONDITION CODE WAS 0
```

Figure 29. An Example of CHKLIST Output

APPENDIX F. COMMAND REFERENCE SUMMARY

This section shows the format of each Access Method Services functional command, parameter abbreviation, and defaults. The commands are grouped together to allow you to remove these pages and use them for a quick reference.

ALTER—MODIFY ATTRIBUTES OF A DATA SET CATALOGED IN A VSAM CATALOG

Parameters	Abbrev.
ALTER	—
entryname[/password]	—
[ADDVOLUMES(volser[b volser...])]	AVOL
[ATTEMPTS(number)]	ATT
[AUTHORIZATION(entrypoint[b string])]	AUTH
[BUFFERSPACE(size)]	BUFSP, BUFSPC
[CODE(code)]	—
[CONTROLPW(password)]	CTLPW
[DESTAGEWAIT NODESTAGEWAIT]	DSTGW NDSTGW
[EMPTY NOEMPTY]	EMP NEMP
[ERASE NOERASE]	ERAS NERAS
[EXCEPTIONEXIT(entrypoint)]	EEXT
[FILE(dname)]	—
[FREESPACE(CI-percent[b CA-percent])]	FSPC
[INHIBIT UNINHIBIT]	INH UNINH
[KEYS(length b offset)]	—
[MASTERPW(password)]	MRPW
[NEWNAME(newname)]	NEWNM
[NULLIFY(NULL
[AUTHORIZATION(MODULE STRING)]	AUTH MDLE
	STRG
	—
[CODE]	CTLPW
[CONTROLPW]	EEXT
[EXCEPTIONEXIT]	MRPW
[MASTERPW]	—
[OWNER]	RDPW
[READPW]	RETN
[RETENTION]	UPDPW
[UPDATEPW]]]	—
[OWNER(ownerid)]	RDPW
[READPW(password)]	RECSZ
[RECORDSIZE(average b maximum)]	RVOL
[REMOVEVOLUMES(volser[b volser...])]	SCR NSCR
[SCRATCH NOSCRATCH]	SHR
[SHAREOPTIONS(crossregion[b crosssystem])]	— — CYLF
[STAGE BIND CYLINDERFAULT]	— —
[TO(date) FOR(days)]	UNQK NUNQK
[UNIQUEKEY NONUNIQUEKEY]	UPD NUPD
[UPDATE NOUPDATE]	UPDPW
[UPDATEPW(password)]	UPG NUPG
[UPGRADE NOUPGRADE]	WCK NWCK
[WRITECHECK NOWRITECHECK]	
[CATALOG(catname[/password])]	CAT

BLDINDEX—BUILD ALTERNATE INDEXES FOR EXISTING VSAM CLUSTERS

Parameters	Abbrev.
BLDINDEX {INFILE(dname[/password]) INDATASET(entryname[/password])} {OUTFILE(dname[/password] [b dname[/password]...]) OUTDATASET(entryname[/password] [b entryname[/password]...])} [CATALOG(catname[/password])] [EXTERNALSORT INTERNALSORT] [WORKFILES(dname b dname)]	BIX IFILE IDS OFILE ODS CAT ESORT ISORT WFILE

CHKLIST—IDENTIFY TAPE VOLUMES MOUNTED WHEN A CHECKPOINT WAS TAKEN

Parameters	Abbrev.
CHKLIST INFILE(dname) [CHECKID(checkid...)] [OUTFILE(dname)]	CKLST IFILE CHKID OFILE

CNVTCAT—CONVERT OS CATALOG ENTRIES INTO OS/VS2 VSAM CATALOG ENTRIES

Parameters	Abbrev.
CNVTCAT {INFILE(dname) INDATASET(entryname)} [CATALOG(catname[/password])] [CVOLEQUATES((catname b(volser [b volser...])][b(catname...)...])] [LIST NOLIST] [MASTERCATALOG(catname[/password])]	CNVTC IFILE IDS CAT CVEQU — NLIST MCAT

DEFINE ALIAS—DEFINE AN ALTERNATE NAME FOR A USER CATALOG OR A NONVSAM DATA SET

Parameters	Abbrev.
DEFINE ALIAS (NAME(aliasname) RELATE(entryname)) [CATALOG(catname[/password])]	DEF — — REL CAT

DEFINE ALTERNATEINDEX—DEFINE A VSAM CATALOG ENTRY FOR AN ALTERNATE INDEX

Parameters	Abbrev.
DEFINE ALTERNATEINDEX	DEF AIX
(NAME(entryname))	—
RELATE(entryname[/password])	REL
{CYLINDERS ¹ (primary[b secondary])}	CYL
RECORDS ¹ (primary[b secondary])}	REC
TRACKS ¹ (primary[b secondary])}	TRK
VOLUMES ² (volser[b volser...])	VOL
[ATTEMPTS(number 2)]	ATT
[AUTHORIZATION(entrypoint[b string])]	AUTH
[BUFFERSPACE(size)]	BUFSP, BUFSPC
[CODE(code)]	—
[CONTROLINTERVALSIZE(size)]	CISZ, CNVSZ
[CONTROLPW(password)]	CTLPW
[DESTAGWAIT NODESTAGWAIT]	DSTGW NDSTGW
[ERASE NOERASE]	ERAS NERAS
[EXCEPTIONEXIT(entrypoint)]	EEXT
[FILE(dname)]	—
[FREESPACE(CI-percent[b CA-percent] 0 0)]	FSPC
[IMBED NOIMBED]	IMBD NIMBD
[KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])]	KRNG
[KEYS(length b offset 64 0)]	—
[MASTERPW(password)]	MRPW
[MODEL(entryname[/password] [b catname[/password]])]	—
[ORDERED UNORDERED]	ORD UNORD
[OWNER(ownerid)]	—
[READPW(password)]	RDPW
[RECORDSIZE(average b maximum 4086 32600)]	RECSZ
[REPLICATE NOREPLICATE]	REPL NREPL
[REUSE NOREUSE]	RUS NRUS
[SHAREOPTIONS(crossregion [b crosssystem] 1 3)]	SHR
[SPEED RECOVERY]	— RCVY
[STAGE BIND CYLINDERFAULT]	— — CYLF
[TO(date) FOR(days)]	— —
[UNIQUE SUBALLOCATION]	UNQ SUBAL
[UNIQUEKEY NONUNIQUEKEY]	UNQK NUNQK
[UPDATEPW(password)]	UPDPW
[UPGRADE NOUPGRADE]	UPG NUPG
[WRITECHECK NOWRITECHECK]	WCK NWCK

(Continued)

DEFINE ALTERNATEINDEX (CONTINUED)

Parameters	Abbrev.
[DATA	—
[[ATTEMPTS(number)]]	ATT
[AUTHORIZATION(entrypoint[b string])]	AUTH
[BUFFERSPACE(size)]]	BUFSP, BUFSPC
[CODE(code)]]	—
[CONTROLINTERVALSIZE(size)]]	CISZ, CNVSZ
[CONTROLPW(password)]]	CTLPW
[CYLINDERS(primary[b secondary])]	CYL
[RECORDS(primary[b secondary])]	REC
[TRACKS(primary[b secondary])]	TRK
[DESTAGWAIT NODESTAGWAIT]	DSTGW NDSTGW
[ERASE NOERASE]	ERAS NERAS
[EXCEPTIONEXIT(entrypoint)]]	EEXT
[FILE(dname)]]	—
[FREESPACE(CI-percent[b CA-percent])]	FSPC
[KEYRANGES((lowkey b highkey)	KRNG
[(lowkey b highkey)...]]]	—
[KEYS(length b offset)]]	—
[MASTERPW(password)]]	MRPW
[MODEL(entryname[/password]	—
[b catname[/password]])]	—
[NAME(entryname)]]	—
[ORDERED UNORDERED]	ORD UNORD
[OWNER(ownerid)]]	—
[READPW(password)]]	RDPW
[RECORDSIZE(average b maximum)]]	RECSZ
[REUSE NOREUSE]	RUS NRUS
[SHAREOPTIONS(crossregion[b crosssystem])]	SHR
[SPEED RECOVERY]	— RCVY
[STAGE BIND CYLINDERFAULT]	— — CYLF
[UNIQUE SUBALLOCATION]	UNQ SUBAL
[UNIQUEKEY NONUNIQUEKEY]	UNQK NUNQK
[UPDATEPW(password)]]	UPDPW
[VOLUMES(volser[b volser...])]	VOL
[WRITECHECK NOWRITECHECK]]]	WCK NWCK

(Continued)

DEFINE ALTERNATEINDEX (CONTINUED)

Parameters	Abbrev.
[INDEX	IX
[[ATTEMPTS(number)]	ATT
[AUTHORIZATION(entrypoint[b string])]	AUTH
[CODE(code)]	—
[CONTROLINTERVALSIZE(size)]	CISZ, CNVSZ
[CONTROLPW(password)]	CTLPW
[CYLINDERS(primary[b secondary])]	CYL
RECORDS(primary[b secondary])]	REC
TRACKS(primary[b secondary])]	TRK
[DESTAGEWAIT NODESTAGEWAIT]	DSTGW NDSTGW
[EXCEPTIONEXIT(entrypoint)]	EEXT
[FILE(dname)]	—
[IMBED NOIMBED]	IMBD NIMBD
[MASTERPW(password)]	MRPW
[MODEL(entryname[/password]	—
[b catname[/password]])]	
[NAME(entryname)]	—
[ORDERED UNORDERED]	ORD UNORD
[OWNER(ownerid)]	—
[READPW(password)]	RDPW
[REPLICATE NOREPLICATE]	REPL NREPL
[REUSE NOREUSE]	RUS NRUS
[SHAREOPTIONS(crossregion[b crosssystem])]	SHR
[STAGE BIND CYLINDERFAULT]	— — CYLF
[UNIQUE SUBALLOCATION]	UNQ SUBAL
[UPDATEPW(password)]	UPDPW
[VOLUMES(volser[b volser...])]	VOL
[WRITECHECK NOWRITECHECK]]]	WCK NWCK
[CATALOG(catname[/password])]	CAT

¹CYLINDERS, RECORDS, or TRACKS must be specified either as a parameter of ALTERNATEINDEX, as a parameter of DATA, or as a parameter of both DATA and INDEX if MODEL is not specified.

²VOLUMES must be specified as a parameter of ALTERNATEINDEX or as a parameter of both DATA and INDEX if MODEL is not specified.

DEFINE CLUSTER—DEFINE A VSAM CATALOG ENTRY FOR A VSAM CLUSTER

Parameters	Abbrev.
DEFINE CLUSTER (NAME(entryname) {CYLINDERS ¹ {primary[b secondary]]} RECORDS ¹ {primary[b secondary]]} TRACKS ¹ {primary[b secondary]]} VOLUMES ² {volser[b volser...]} [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [DESTAGWAIT NODESTAGWAIT] [ERASE NOERASE] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE ³ {CI-percent [b CA-percent] 0 0}] [IMBED NOIMBED] ³ [INDEXED NONINDEXED NUMBERED] [KEYRANGES ³ {(lowkey b highkey) [(lowkey b highkey)...]}] [KEYS ³ {length b offset 64 0}] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [ORDERED UNORDERED] [OWNER(ownerid)] [READPW(password)] [RECORDSIZE ⁴ {average b maximum}] [REPLICATE NOREPLICATE] ³ [REUSE NOREUSE] [SHAREOPTIONS(crossregion [b crosssystem] 1 3)] [SPANNED NONSPANNED] [SPEED RECOVERY] [STAGE BIND CYLINDERFAULT] [TO(date) FOR(days)] [UNIQUE SUBALLOCATION] [UPDATEPW(password)] [WRITECHECK NOWRITECHECK])	DEF CL — CYL REC TRK VOL ATT AUTH BUFSP, BUFSPC — CISZ, CNVSZ CTLPW DSTGW NDSTGW ERAS NERAS EEXT — FSPC — IMBD NIMBD IXD NIXD NUMD KRNG — MRPW — ORD UNORD — RDPW RECSZ REPL NREPL RUS NRUS SHR — SPND NSPND — RCVY — — CYLF — — UNQ SUBAL UPDPW WCK NWCK — ATT AUTH BUFSP, BUFSPC — CISZ, CNVSZ CTLPW CYL REC TRK DSTGW NDSTGW ERAS NERAS EEXT — FSPC KRNG — MRPW
[DATA {[ATTEMPTS(number)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size)] [CODE(code)] [CONTROLINTERVALSIZE(size)] [CONTROLPW(password)] [CYLINDERS(primary[b secondary]] RECORDS(primary[b secondary]] TRACKS(primary[b secondary]))] [DESTAGWAIT NODESTAGWAIT] [ERASE NOERASE] [EXCEPTIONEXIT(entrypoint)] [FILE(dname)] [FREESPACE ³ {CI-percent[b CA-percent]]] [KEYRANGES ³ {(lowkey b highkey) [(lowkey b highkey)...]}] [KEYS ³ {length b offset}] [MASTERPW(password)]	— ATT AUTH BUFSP, BUFSPC — CISZ, CNVSZ CTLPW CYL REC TRK DSTGW NDSTGW ERAS NERAS EEXT — FSPC KRNG — MRPW

(Continued)

DEFINE CLUSTER (CONTINUED)

Parameters	Abbrev.
[MODEL(entryname[/password] [b catname[/password]])]	—
[NAME(entryname)]	—
[ORDERED UNORDERED]	ORD UNORD
[OWNER(ownerid)]	—
[READPW(password)]	RDPW
[RECORDSIZE(average b maximum)]	RECSZ
[REUSE NOREUSE]	RUS NRUS
[SHAREOPTIONS(crossregion[b crosssystem])]	SHR
[SPANNED NONSPANNED]	SPND NSPND
[SPEED RECOVERY]	— RCVY
[STAGE BIND CYLINDERFAULT]	— — CYLF
[UNIQUE SUBALLOCATION]	UNQ SUBAL
[UPDATEPW(password)]	UPDPW
[VOLUMES(volser[b volser...])]	VOL
[WRITECHECK NOWRITECHECK]]	WCK NWCK
[INDEX ³	IX
([ATTEMPTS(number)]	ATT
[AUTHORIZATION(entrypoint[b string]])]	AUTH
[CODE(code)]	—
[CONTROLINTERVALSIZE(size)]	CISZ,CNVSZ
[CONTROLPW(password)]	CTLPW
[CYLINDERS(primary[b secondary])]	CYL
RECORDS(primary[b secondary])]	REC
TRACKS(primary[b secondary])]	TRK
[DESTAGEWAIT NODESTAGEWAIT]	DSTGW NDSTGW
[EXCEPTIONEXIT(entrypoint)]	EEXT
[FILE(dname)]	—
[IMBED NOIMBED] ³	IMBD NIMBD
[MASTERPW(password)]	MRPW
[MODEL(entryname[/password] [b catname[/password]])]	—
[NAME(entryname)]	—
[ORDERED UNORDERED]	ORD UNORD
[OWNER(ownerid)]	—
[READPW(password)]	RDPW
[REPLICATE NOREPLICATE] ³	REPL NREPL
[REUSE NOREUSE]	RUS NRUS
[SHAREOPTIONS(crossregion[b crosssystem])]	SHR
[STAGE BIND CYLINDERFAULT]	— — CYLF
[UNIQUE SUBALLOCATION]	UNQ SUBAL
[UPDATEPW(password)]	UPDPW
[VOLUMES(volser[b volser...])]	VOL
[WRITECHECK NOWRITECHECK]]	WCK NWCK
[CATALOG(catname[/password]])]	CAT

¹CYLINDERS, RECORDS, or TRACKS must be specified either as a parameter of CLUSTER, as a parameter of DATA, or as a parameter of both DATA and INDEX if MODEL

²VOLUMES must be specified as a parameter of CLUSTER or as a parameter of both DATA and INDEX if MODEL is not specified.

³Can only be specified for a key-sequenced cluster.

⁴When SPANNED is specified, the default is RECORDSIZE (4086 32600). Otherwise, the default is RECORDSIZE (4089 4089).

DEFINE GENERATIONDATAGROUP—CREATE A CATALOG ENTRY FOR A GENERATION DATA GROUP

Parameters	Abbrev.
DEFINE GENERATIONDATAGROUP (NAME(entryname) LIMIT(limit) [EMPTY NOEMPTY] [OWNER(ownerid)] [SCRATCH NOSCRATCH] [TO(date) FOR(days)]) [CATALOG(catname[/password])]	DEF GDG — LIM EMP NEMP — SCR NSCR — CAT

DEFINE NONVSAM—DEFINE A VSAM CATALOG ENTRY FOR A NONVSAM DATA SET

Parameters	Abbrev.
DEFINE NONVSAM (NAME(entryname) DEVICETYPES(devtype[b devtype...]) VOLUMES(volser[b volser...]) [FILESEQUENCENUMBERS(number [b number...])] [OWNER(ownerid)] [TO(date) FOR(days)]) [CATALOG(catname[/password])]	DEF NVSAM — DEVT VOL FSEQN — — CAT

DEFINE PAGESPACE—DEFINE A VSAM CATALOG ENTRY FOR A PAGE SPACE

Parameters	Abbrev.
DEFINE PAGESPACE (NAME(entryname) (CYLINDERS(primary) RECORDS(primary) TRACKS(primary)) VOLUME(volser) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLPW(password)] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [SWAP NOSWAP] [TO(date) FOR(days)] [UNIQUE SUBALLOCATION] [UPDATEPW(password)]) [CATALOG(catname[/password])]	DEF PGSPC — CYL REC TRK VOL ATT AUTH — CTLPW — MRPW — RDPW — NSWAP — UNQ SUBAL UPDPW CAT

DEFINE PATH—DEFINE A RELATIONSHIP BETWEEN AN ALTERNATE INDEX AND ITS BASE CLUSTER

Parameters	Abbrev.
DEFINE PATH (NAME(entryname) PATHENTRY(entryname[/password]) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [CODE(code)] [CONTROLPW(password)] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [TO(date) FOR(days)] [UPDATE NOUPDATE] [UPDATEPW(password)] [CATALOG(catname[/password])]	DEF — — PENT ATT AUTH — CTLPW — MRPW — — RDPW — — UPD NUPD UPDPW CAT

DEFINE SPACE—DEFINE A VSAM DATA SPACE

Parameters	Abbrev.
DEFINE SPACE ({CANDIDATE CYLINDERS(primary[b secondary]) RECORDS(primary[b secondary] b RECORDSIZE(average b maximum) TRACKS(primary[b secondary]) VOLUMES(volser[b volser...]) [FILE(dname)]) [CATALOG(catname[/password])]	DEF SPC CAN CYL REC RECSZ TRK VOL — CAT

DEFINE USERCATALOG—CREATE A VSAM USER CATALOG

Parameters	Abbrev.
DEFINE USERCATALOG MASTERCATALOG (NAME(entryname) {CYLINDERS(primary[b secondary]) RECORDS(primary[b secondary]) TRACKS(primary[b secondary])} VOLUME(volser) [ATTEMPTS(number 2)] [AUTHORIZATION(entrypoint[b string])] [BUFFERSPACE(size 3072)] [CODE(code)] [CONTROLPW(password)] [DESTAGWAIT <u>NODESTAGWAIT</u>] [FILE(dname)] [MASTERPW(password)] [MODEL(entryname[/password] [b catname[/password]])] [OWNER(ownerid)] [READPW(password)] [RECOVERABLE <u>NOTRECOVERABLE</u>] [TO(date) FOR(days)] [UPDATEPW(password)] [WRITECHECK <u>NOWRITECHECK</u>])	DEF UCAT MCAT — CYL REC TRK VOL ATT AUTH BUFSP, BUFSPC — CTLPW DSTGW NDSTGW — MRPW — RDPW RVBL NRVBL — UPDPW WCK NWCK —
[DATA ([BUFFERSPACE(size)] [CYLINDERS(primary[b secondary]) RECORDS(primary[b secondary]) TRACKS(primary[b secondary])] [DESTAGWAIT NODESTAGWAIT] [RECOVERABLE NOTRECOVERABLE] [WRITECHECK NOWRITECHECK])	BUFSP, BUFSPC CYL REC TRK DSTGW NDSTGW RVBL NRVBL WCK NWCK
[INDEX ([CYLINDERS(primary) RECORDS(primary) TRACKS(primary)] [DESTAGWAIT NODESTAGWAIT] [WRITECHECK NOWRITECHECK])	IX CYL REC TRK DSTGW NDSTGW WCK NWCK
[CATALOG(mastercatname[/password])]	CAT

DELETE—DELETE ENTRIES FROM A VSAM CATALOG

Parameters	Abbrev.
DELETE {entryname[/password] [b entryname[/password]...]} [ALIAS ALTERNATEINDEX CLUSTER GENERATIONDATAGROUP NONVSAM PAGESPACE PATH SPACE ¹ USERCATALOG ¹] [ERASE NOERASE] [FILE(dname)] [FORCE NOFORCE] [PURGE NOPURGE] [SCRATCH NOSCRATCH] [CATALOG(catname[/password])]	DEL — — AIX CL GDG NVSAM PGSPC — SPC UCAT ERAS NERAS — FRC NFRC PRG NPRG SCR NSCR CAT

¹When you delete a data space or catalog, you cannot delete any other type of entry. You must identify the type of entry to be deleted, by specifying SPACE or USERCATALOG.

EXPORT—DISCONNECT A VSAM USER CATALOG

Parameters	Abbrev.
EXPORT usercatname[/password] DISCONNECT	EXP — DCON

EXPORT—COPY A VSAM CLUSTER OR ALTERNATE INDEX, TOGETHER WITH ITS CATALOG ENTRIES, TO A SEQUENTIAL DATA SET

Parameters	Abbrev.
EXPORT entryname[/password] {OUTFILE(dname) OUTDATASET(entryname)} [INFILE(dname)] [ERASE NOERASE] [INHIBITSOURCE NOINHIBITSOURCE] [INHIBITTARGET NOINHIBITTARGET] [PURGE NOPURGE] [TEMPORARY PERMANENT]	EXP — OFILE ODS IFILE ERAS NERAS INHS NINHS INHT NINHT PRG NPRG TEMP PERM

EXPORTRA—COPY CATALOG ENTRIES FROM THE CATALOG RECOVERY AREA (CRA)

Parameters	Abbrev.
EXPORTRA OUTFILE(dname) CRA((dname1 {ALL[b INFILE(dname2)]} ENTRIES((entryname[b dname3]) [(entryname[b dname3])...]) NONE) [(dname1...)] [FORCE NOFORCE] [MASTERPW(password)]	XPRA OFILE — — IFILE ENT — FRC NFRC MRPW

IMPORT—CONNECT A VSAM USER CATALOG

Parameters	Abbrev.
IMPORT CONNECT OBJECTS((usercatname DEVICETYPE(devtype) VOLUMES(volser))) [CATALOG(mastercatname[/password])]	IMP CON OBJ DEVT VOL CAT

IMPORT—RESTORE A VSAM CLUSTER OR ALTERNATE INDEX FROM A PORTABLE DATA SET CREATED BY THE EXPORT COMMAND

Parameters	Abbrev.
IMPORT {INFILE(dname) INDATASET(entryname)} {OUTFILE(dname[/password]) OUTDATASET(entryname[/password])} [ERASE NOERASE] [INTOEMPTY] [OBJECTS((entryname [FILE(dname)] [KEYRANGES((lowkey b highkey) [(lowkey b highkey)...])] [NEWNAME(newname)] [ORDERED UNORDERED] [VOLUMES(volser[b volser...])] [(entryname...)...])] [PURGE NOPURGE] [SAVRAC NOSAVRAC] [CATALOG(catname[/password])]	IMP IFILE IDS OFILE ODS ERAS NERAS IEMPTY OBJ — KRNG NEWNM ORD UNORD VOL PRG NPRG — — CAT

IMPORTRA—RESTORE CATALOG ENTRIES FROM A PORTABLE DATA SET CREATED BY THE EXPORTRA COMMAND

Parameters	Abbrev.
IMPORTRA {INFILE(dname) INDATASET(entryname)} [OUTFILE(dname)] [OBJECTS(entryname [DEVICETYPES(devtype[b devtype...])] [FILE(dname)] [VOLUMES(volser[b volser...])] [(entryname...)...])] [SAVRAC NOSAVRAC] [CATALOG(catname[/password])]	MPRA IFILE IDS OFILE OBJ DEVT — VOL — — CAT

LISTCAT—LIST ENTRIES FROM A CATALOG

Parameters	Abbrev.
LISTCAT [CATALOG(catname[/password])] [ENTRIES(entryname[/password] [b entryname[/password]...])] [LEVEL(level)] [ALIAS] [b ALTERNATEINDEX] [b CLUSTER] [b DATA] [b GENERATIONDATAGROUP] [b INDEX] [b NONVSAM] [b PAGESPACE] [b PATH] [b SPACE] [b USERCATALOG] [ALL] ALLOCATION HISTORY NAME VOLUME [CREATION(days)] [EXPIRATION(days)] [NOTUSABLE] [OUTFILE(dname)]	LISTC CAT ENT LVL — AIX CL — GDG IX NVSAM PGSPC — SPC UCAT — ALLOC HIST — VOL CREAT EXPIR NUS OFFILE

LISTCRA—LIST THE CONTENTS OF THE CATALOG RECOVERY AREA (CRA)

Parameters	Abbrev.
LISTCRA INFILE(dname[b dname...]) [CATALOG(catname[/password]b dname)] [COMPARE NOCOMPARE] [DUMP NAME SEQUENTIALDUMP] [MASTERPW(password)] [OUTFILE(dname)]	LISTR IFILE CAT CMPR NCMPR — — SDUMP MRPW OFFILE

PRINT—PRINT THE CONTENTS OF A DATA SET

Parameters	Abbrev.
PRINT {INFILE(dname[/password]) INDATASET(entryname[/password])} [OUTFILE(dname)] [CHARACTER DUMP HEX] [FROMKEY(key) FROMADDRESS(address)] FROMNUMBER(number) SKIP(count)] [TOKEY(key) TOADDRESS(address)] TONUMBER(number) COUNT(count)]	IFILE IDS OFILE CHAR — — FKEY FADDR FNUM — — TADDR TNUM —

REPRO—COPY DATA SETS

Parameters	Abbrev.
REPRO {INFILE(dname[/password] [b ENVIRONMENT(DUMMY)]) INDATASET(entryname[/password] [b ENVIRONMENT(DUMMY)])} {OUTFILE(dname[/password]) OUTDATASET(entryname[/password])} [REPLACE NOREPLACE] [REUSE NOREUSE] [FROMKEY(key) FROMADDRESS(address)] FROMNUMBER(number) SKIP(count)] [TOKEY(key) TOADDRESS(address)] TONUMBER(number) COUNT(count)]	— IFILE ENV DUM IDS ENV DUM OFILE ODS REP NREP RUS NRUS FKEY FADDR FNUM — — TADDR TNUM —

RESETCAT—SYNCHRONIZE A RECOVERABLE CATALOG TO THE LEVEL OF ITS OWNED VOLUMES

Parameters	Abbrev.
RESETCAT CATALOG(catname[/password][b dname]) {CRAFILES((dname1[b ALL b NONE]) [(dname1[b ALL b NONE])...]) CRAVOLUMES((volser[b devtype]) [(volser[b devtype])...])} [IGNORE NOIGNORE] [MASTERPW(password)] [WORKCAT(catname[/password])] [WORKFILE(dname[/password])]	RCAT CAT — — — — — CRAVOL IGN NIGN MRPW WCAT WFILE

VERIFY—RESTORE A VSAM CLUSTER'S END-OF-FILE VALUES

Parameters	Abbrev.
VERIFY {FILE(dname[/password]) DATASET(entryname[/password])}	VFY — DS

EXECUTION CONTROL COMMANDS

IF-THEN-ELSE COMMAND SEQUENCE

Parameters	Abbrev.
IF {LASTCC MAXCC}[b]comparand[b]number THEN[b command] DO command set END] [ELSE[b command] DO command set END]]	

PARAM COMMAND

Parameters	Abbrev.
PARAM [TEST([[TRACE] [AREAS(areaid[b areaid...]]) [FULL((dumpid[b count1[b count2]]) [(dumpid...)...]])] OFF]]) [GRAPHICS(CHAIN(chain) TABLE(mname))] [MARGINS(leftmargin b rightmargin)]	

SET COMMAND

Parameters	Abbrev.
SET {MAXCC LASTCC}[b]=[b]number	

APPENDIX G. INVOKING ACCESS METHOD SERVICES FROM A PROBLEM PROGRAM

Access Method Services can be invoked by a problem program through the use of the ATTACH, LINK, or LOAD and CALL macro instructions.

The dynamic invocation of Access Method Services enables respecification of selected processor defaults as well as the ability to manage input/output operations for selected data sets.

AUTHORIZED PROGRAM FACILITY (APF)

There are six Access Method Services commands which require APF authorization; they are:

- REPRO—when using the copy-catalog or catalog unload/reload facility.
- PRINT—when the object to be printed is a VSAM catalog.
- LISTCRA
- EXPORTRA
- RESETCAT
- VERIFY—when a VSAM catalog is to be verified.

If the problem program invoking Access Method Services needs to perform any of these functions, it must be authorized and be located in an authorized library. Additionally, a User I/O Routine must reside in an authorized library if the problem program it is associated with is authorized. See "Authorized Program Facility (APF)" in OS/VS2 System Programming Library: Supervisor, for information about program authorization.

INVOKING MACRO INSTRUCTIONS

The following descriptions of the invoking macro instructions are related to Figure 30, which describes the argument lists referenced by the invoking macros.

LINK OR ATTACH MACRO INSTRUCTION

Access Method Services may be invoked through either the LINK or the ATTACH macro instruction.

The format of the LINK or ATTACH macro instruction is:

```
[name]
  {LINK|ATTACH}
  EP=IDCAMS,
  PARAM={optionaddr
        [,dnameaddr]
        [,pgnoaddr]
        [,iolistaddr]},
  VL=1
```

EP=IDCAMS
specifies that the program to be invoked is IDCAMS.

PARAM=

specifies the addresses of the parameters to be passed to IDCAMS. These values can be coded:

optionaddr

specifies the address of an option list, which can be specified in the PARM parameter of the EXEC statement and is a valid set of parameters for the Access Method Services PARM command. If you do not wish to specify any options, this address must point to a halfword of binary zeros. Figure 30 shows the format of the options list.

dnameaddr

specifies the address of a list of alternate DD names for standard data sets used during IDCAMS processing. If standard DD names are used and this is not the last parameter in the list, it should point to a halfword of binary zeros. If it is the last parameter, it may be omitted. Figure 30 shows the format of the alternate DD name list.

pgnoaddr

specifies the address of a 6-byte area which contains an EBCDIC starting page number for the system output file. If the page number is not specified, but this is not the last parameter in the list, the parameter must point to a halfword of binary zeros. If it is the last parameter, it may be omitted. If omitted the default page number is 1. Figure 30 shows the format of the page-number area.

iolistaddr

specifies the address of a list of externally-controlled data sets and the addresses of corresponding I/O routines. If no external I/O routines are supplied, this parameter may be omitted. Figure 30 shows the format of the I/O list.

VL=1

causes the high-order bit of the last address parameter of the PARAM list to be set to 1.

LOAD AND CALL MACRO INSTRUCTIONS

Access Method Services may also be invoked via a LOAD of the module IDCAMS, followed by a CALL to that module. The format of the LOAD macro instruction is:

[name]	LOAD	{EP=IDCAMS EPLOC=address of name}
--------	------	-----------------------------------

where:

EP=IDCAMS

is the entry point name of the IDCAMS program to be loaded into virtual storage.

EPLOC=address of name

is the address of an 8-byte character string 'IDCAMS bb'.

After loading IDCAMS, Register 15 must be loaded with the address returned from the LOAD macro. Then CALL may be used to pass control to IDCAMS. The format of the CALL macro instruction is:

[name]	LR CALL	15,0 (15), (optionaddr [,dnameaddr] [,pgnoaddr] [,iolistaddr]), VL
--------	------------	--

where:

15 is the register containing the address of the entry point to be given control.

optionaddr specifies the address of an options list which can be specified in the PARM parameter of the EXEC statement and is a valid set of parameters for the Access Method Services PARM command. If you do not wish to specify any options, this address must point to a halfword of binary zeros. Figure 30 shows the format of the options list.

dnameaddr specifies the address of a list of alternate DD names for standard data sets used during IDCAMS processing. If standard DD names are used and this is not the last parameter in the list, it should point to a halfword of binary zeros. If it is the last parameter, it may be omitted. Figure 30 shows the format of the alternate DD name list.

pgnoaddr specifies the address of a 6-byte area which contains an EBCDIC starting page number for the system output file. If the page number is not specified, but this is not the last parameter in the list, the parameter must point to a halfword of binary zeros. If it is the last parameter, it may be omitted. If omitted the default page number is 1. Figure 30 shows the format of the page number area.

iolistaddr specifies the address of a list of externally controlled data sets and the addresses of corresponding I/O routines. If no external I/O routines are supplied, this parameter may be omitted. Figure 30 shows the format of the I/O list.

VL causes the high-order bit of the last address parameter in the macro expansion to be set to 1.

INVOCATION FROM A PL/I PROGRAM

Access Method Services may also be invoked from a PL/I program using the facilities of the IBM PL/I Optimizing Compiler Program Product. IDCAMS must be declared to the compiler as an external entry point with the ASSEMBLER and INTER options. The Access Method Services processor is loaded by issuing a FETCH IDCAMS statement, is branched to via a CALL statement, and deleted via a RELEASE IDCAMS statement. The format of the CALL statement is:

CALL	IDCAM	(options[,dnames][,pageno][,iolist]);
------	-------	---------------------------------------

where:

options

specifies a valid set of parameters for the Access Method Services PARM command. If no parameters are to be specified, options should be a halfword of binary zeros. Figure 30 shows the format of the options area.

dnames

specifies a list of alternate DD names for standard data sets used during IDCAMS processing. If standard DD names are used and this is not the last parameter in the list, dnames should be a halfword of binary zeros. If it is the last parameter, it may be omitted. Figure 30 shows the format of the alternate DD names list.

pageno

specifies a 6-byte field that contains an EBCDIC starting page number for the system output file. If the page number is not specified, but this is not the last parameter in the list, the parameter must be a halfword of binary zeros. If it is the last parameter, it may be omitted. If not specified, the default page number is 1. Figure 30 shows the format of the page number area.

iolist

specifies a list of externally controlled data sets and the addresses of corresponding I/O routines. If no external I/O routines are supplied, this parameter may be omitted. Figure 30 shows the format of the I/O list.

PROCESSOR INVOCATION

Figure 30 shows the processor invocation argument list as it exists in the user's area.

Entry and exit to the Access Method Services processor occur through IDCSA01, a module of the System Adapter. Standard linkage is used; that is, register 1 points to the argument list, register 13 points to a save area, register 14 contains the return address, and register 15 contains the entry point address for IDCSA01. On exit from the Access Method Services processor, register 15 contains the value of MAXCC (see the section "Processor Condition Codes" below).

The argument list, as shown in Figure 30, can be a maximum of four fullword addresses pointing to strings of data. The last address in the list contains a "1" in the sign field. The first three possible strings of data each begin with a two-byte length field. A null element in the list can be indicated by either an address of zeros or a length of zero.

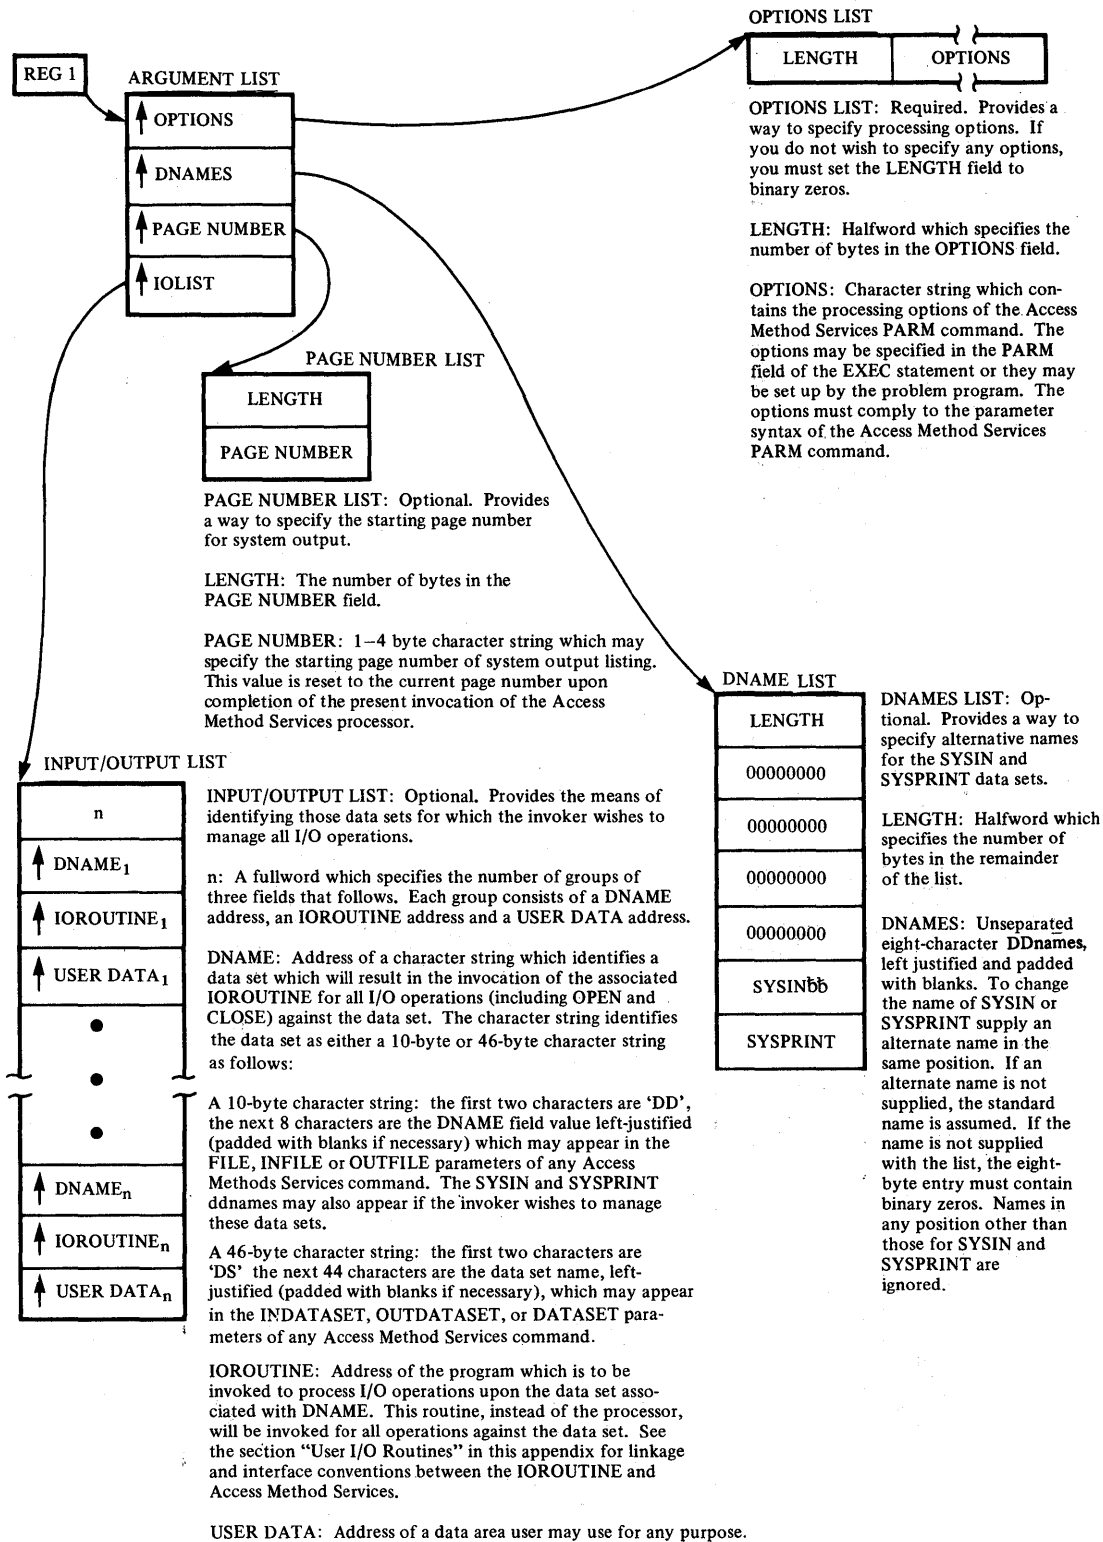


Figure 30. Processor Invocation Argument List From a Problem Program

PROCESSOR CONDITION CODES

The processor's condition code is LASTCC, which can be interrogated in the command stream following each functional command. The possible values, their meanings, and examples of causes are in the following table. The table illustrates the value of LASTCC.

Code Meaning

0	The function was executed successfully. Informational messages may have been issued.
4	Some annoyance in executing the complete function was met, but it was possible to continue. The results may not be exactly what the user wants, but no permanent harm appears to have been done by continuing. A warning message was issued.
8	A function could not perform all that was asked of it. The function was completed, but specific details were bypassed.
12	The entire function could not be performed.
16	Severe error or problem encountered. Remainder of command stream is flushed and processor returns condition code 16 to the operating system.

LASTCC is set by the processor at the completion of each functional command. MAXCC, which can also be interrogated in the command stream, is the highest value of LASTCC thus far encountered.

USER I/O ROUTINES

User I/O routines enable a user to perform all I/O operations for a data set which would normally be handled by the Access Methods Services processor. This makes it possible, for instance, to control the command input stream by providing an I/O routine for SYSIN.

A user I/O routine is invoked by Access Method Services for all operations against the selected data sets. The identification of the data sets and their associated I/O routines is via the input/output list of the processor invocation parameter list (see Figure 30).

When writing a user I/O routine, the user must be aware of three things: First, the processor handles the user data set as if it were a nonVSAM data set that contains variable-length unblocked records (maximum record length is 32,760 bytes) with a physical sequential organization. The processor does not test for the existence of the data set. Second, the user must know the data format so that the user's routine can be coded to handle the correct type of input and format the correct type of output. Third, each user routine must handle errors encountered for data sets it is managing and provide a return code to the processor in register 15. The processor uses the return code to determine what it is to do next.

The permissible return codes are:

- 0—operation successful
- 4—end of data for a GET operation
- 8—error encountered during a GET/PUT operation, but continue processing

- 12—do not allow any further calls (except CLOSE) to this routine

Figure 31 shows the argument list used in communication between the user I/O routine and the Access Method Services processor. The User I/O routine is invoked by the processor for OPEN, CLOSE, GET, and PUT routines.

The type of operation to be performed is indicated via IOFLAGS. The IOINFO field indicates, for OPEN and CLOSE operations, the data set name or ddname of the data set; for GET and PUT operations, the IOINFO field is used to communicate the record length and address.

A user I/O routine for SYSPRINT receives control each time the processor issues a PUT against the SYSPRINT data set. If the PUT has been issued to print an IDC message, the unique message number is passed to the routine via IOFLAGS (see Figure 31). Each IDC message is in the form IDCsnnnI or IDCsnnnnI, where:

s is a code indicating the severity of the problem

nnn or nnnn is the message number which is unique across all IDC messages

The 2-byte message number passed via IOFLAGS is the nnn or nnnn portion of the message converted to binary.

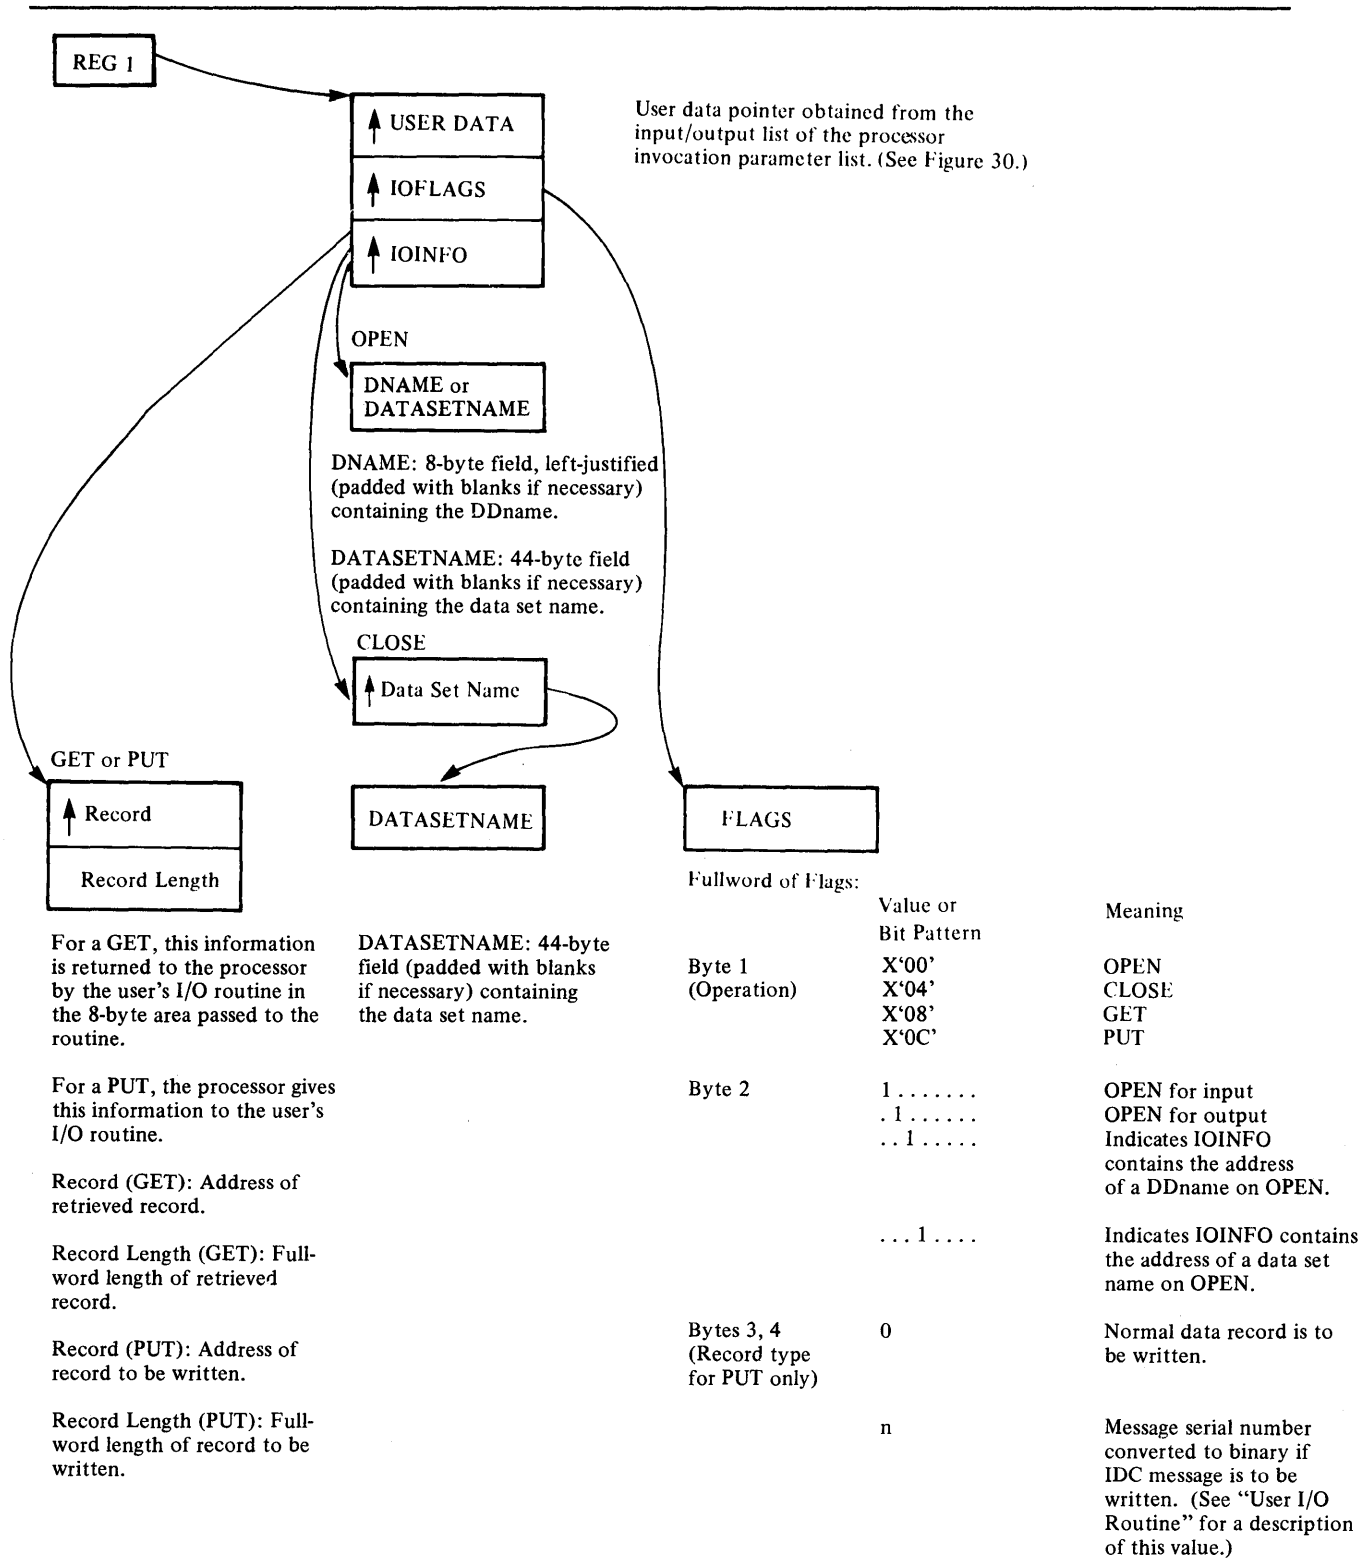


Figure 31. Arguments Passed to and From a User I/O Routine

APPENDIX H. CHANGING THE ATTRIBUTES OF THE VSAM MASTER CATALOG

At some point in time in an existing OS/VS2 MVS system, you may wish to change some of the attributes of your VSAM master catalog. For example, you may wish to convert the VSAM master catalog from nonrecoverable to recoverable, or take advantage of a new VSAM enhancement. This is an example of the conversion process. A description of the process precedes an actual example of each step. When you understand the example's objectives, you can modify it so that it is appropriate for your system.

1. Install on your system the level (or higher level) of VSAM and Access Method Services that supports the features you desire.
2. Export (with the PERMANENT attribute) all VSAM data sets from the old master catalog. Ignore for the moment the storage index and page data sets in the old master catalog.
3. Mount a scratch pack to hold a temporary catalog, a storage index, and the page data sets.
4. Define a temporary master catalog on the scratch pack. You can specify either DEFINE MASTERCATALOG or DEFINE USERCATALOG. In either case, VSAM creates a user catalog which can then be referenced in step 6 and used for the IPL in step 7.
5. Catalog the nonVSAM data sets needed for an IPL on the temporary master catalog. Define a new storage index and page data sets on the scratch volume in the temporary catalog. The page data sets should have the same name as the original page data sets so that the IEASYS00 member in SYS1.PARMLIB need not be replaced.
6. Alter the SYSCATLG member of the SYS1.NUCLEUS data set so that it points to the temporary catalog.
7. IPL the system (cold start procedure). The system now is using the temporary catalog as the master catalog, and refers to the storage index data set and page data sets cataloged in the temporary catalog.
8. Use the ALTER REMOVEVOLUMES command to remove VSAM from the volumes containing the old master catalog, the old page data sets, and the old storage index, and to remove VSAM from any volumes owned by the old master catalog.
9. Define the new catalog in place of the original. This catalog will eventually become the VSAM master catalog. As in step 4, you can specify either DEFINE MASTERCATALOG or DEFINE USERCATALOG.
10. Redefine the storage index and page data sets in their original locations on the new master catalog. (Allow room for the catalog recovery area if the new catalog has been defined with the recoverable attribute.) Catalog all nonVSAM data sets needed for IPL in the new master catalog.
11. Replace the SYSCATLG member in SYS1.NUCLEUS (now cataloged in the new master catalog) so that the SYSCATLG member points to the new master catalog.
12. ReIPL the system (a cold start). The system should now be using the new VSAM master catalog, the redefined storage index, and the newly defined page data sets.
13. Clean up the scratch pack using the ALTER command with the REMOVEVOLUMES parameter.

14. Recatalog all the remaining nonVSAM data sets into the new master catalog.
15. Connect (and define aliases for) OS catalogs (CVOLs) that were connected to the old master catalog.
16. Redefine the data spaces that were deleted from the old master catalog.
17. Import the previously exported VSAM data sets into the new master catalog.
18. Import (with the CONNECT parameter) and define aliases for user catalogs that were connected to the old master catalog.

Your new master catalog is now reestablished. In this example, the new master catalog is recoverable (that is, its volume includes a catalog recovery area and can be processed with the LISTCRA, EXPORTRA, IMPORTRA, and RESETCAT commands).

Figure 32 shows an example of the job stream needed to make your master catalog a recoverable catalog. The step numbers in the jobs correspond to the steps described previously.

If your old master catalog is already recoverable, you can use the EXPORTRA and IMPORTRA commands rather than the EXPORT and IMPORT commands as used in this example. You can use EXPORTRA for all volumes owned by the master catalog, including the master catalog volume, in place of step 2 above. After redefining any VSAM data spaces (step 16), you can use IMPORTRA in place of steps 14, 15, 17, and 18.

```

//JOB1 JOB ...
//* STEP 1 - INSTALL VSAM ON YOUR SYSTEM
//STEP2 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//* DD statements that describe the output data sets
//SYSIN DD *
EXPORT commands required to export the VSAM data sets and user
catalog connector entries from the old master catalog
/*
//* STEP 3 - MOUNT A SCRATCH PACK
//STEP4 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CATPAK DD UNIT=3330,VOL=SER=CATPAK,DISP=OLD
//SYSIN DD *
DEFINE MASTERCATALOG -
(NAME(TEMPCAT) -
FILE(CATPAK) -
VOLUME(CATPAK) -
CYLINDERS (12 1))
/*
//STEP5 EXEC PGM=IDCAMS,COND=(0,LT)
//STEPCAT DD DSN=TEMPCAT,DISP=OLD
//SYSPRINT DD SYSOUT=A
//CATPAK DD UNIT=3330,VOL=SER=CATPAK,DISP=OLD
//SYSIN DD *
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.LPALIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.LINKLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.DSSVM ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.SVCLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.NUCLEUS ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.DCMLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.INDMAC ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.CMDLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.CMDLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.HELP ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.SAMPLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.MACLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.PROCLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.TELCMLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.UADS ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.VTAMLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.IMAGELIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.PARMLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.BROADCAST ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.MANX ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.MANY ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.DUMP00 ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.DUMP01 ) VOL(SG2001) DEVT(3330))
DEF CAT(TEMPCAT) NVSAM (NAME(SYS1.LOGREC ) VOL(SG2001) DEVT(3330))
DEF CLUSTER -
(NAME(SYS1.STGINDEX) -
FILE (CATPAK) -
KEYS(12,8) -
CYLINDERS (6,0) -
RECORDSIZE(2041,2041) -
VOLUME(CATPAK) -
UNIQUE) -
DATA(CONTROLINTERVALSIZE(2048)) -
CATALOG(TEMPCAT)

```

Figure 32 (Part 1 of 4). Sample Job Stream to Make the Master Catalog Recoverable

```

DEFINE PAGESPACE -
  (NAME(SPOOL1P1) -
  FILE(CATPAK) -
  CYLINDERS(80) -
  VOLUME(CATPAK) -
  UNIQUE) -
  CATALOG(TEMPCAT)
DEFINE PAGESPACE -
  (NAME(SPOOL1P2) -
  FILE(CATPAK) -
  CYLINDERS(40) -
  VOLUME(CATPAK) -
  UNIQUE) -
  CATALOG(TEMPCAT)
/*
//STEP6A EXEC PGM=IEHPROGM,COND=(0,LT)
//SG2001 DD DISP=OLD,VOL=SER=SG2001,UNIT=3330
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
SCRATCH DSN=SYS1.NUCLEUS,VOL=3330=SG2001,MEMBER=SYSCATLG
/*
//STEP6B EXEC PGM=IEBGENER,COND=(0,LT)
//SYSIN DD DUMMY
//SYSPRINT DD SYSOUT=A
//SYSUT2 DD DSN=SYS1.NUCLEUS(SYSCATLG),DISP=(OLD,KEEP),
// DCB=BLKSIZE=13030,VOL=SER=SG2001,UNIT=3330
//SYSUT1 DD *
CAPAK (See Note 1)
/*
//* STEP 7 - IPL THE SYSTEM
//
//JOB2 JOB ...
//STEP8 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//DD3330 DD UNIT=3330,DISP=OLD,VOL=SER=(SG2001,SPOOL1)
//SYSIN DD *
ALTER TEMPCAT RVOL ( SG2001 SPOOL1) FILE ( DD3330 )
/*
//STEP9 EXEC PGM=IDCAMS,COND=(0,LT)
//SYSPRINT DD SYSOUT=A
//SG2001 DD UNIT=3330,VOL=SER=SG2001,DISP=OLD
//SYSIN DD *

```

Figure 32 (Part 2 of 4). Sample Job Stream to Make the Master Catalog Recoverable

```

DEFINE MASTERCATALOG -
  (NAME(AMASTCAT) -
  FILE(SG2001) -
  VOLUME(SG2001) -
  CYLINDERS(12 1) -
  RECOVERABLE)
/*
//STEP10 EXEC PGM=IDCAMS,COND=(0,LT)
//STEPCAT DD DSN=AMASTCAT,DISP=OLD
//SYSPRINT DD SYSOUT=A
//SPOOL1 DD UNIT=3330,VOL=SER=SP00L1,DISP=OLD
//SG2001 DD UNIT=3330,VOL=SER=SG2001,DISP=OLD
//SYSIN DD *
  DEF CLUSTER -
    (NAME(SYS1.STGINDEX) -
    FILE (SG2001) -
    KEYS(12,8) -
    CYLINDERS(6,0) -
    RECORDSIZE(2041,2041) -
    VOLUME(SG2001) -
    UNIQUE) -
    DATA(CONTROLINTERVALSIZE (2048)) -
    CATALOG(AMASTCAT)
  DEFINE PAGESPACE -
    (NAME(SPOOL1P1) -
    FILE(SPOOL1) -
    CYLINDERS(140) -
    VOLUME(SPOOL1) -
    UNIQUE) -
    CATALOG(AMASTCAT)
  DEFINE PAGESPACE -
    (NAME(SPOOL1P2) -
    FILE(SPOOL1) -
    CYLINDERS(40) -
    VOLUME(SPOOL1) -
    UNIQUE) -
    CATALOG(AMASTCAT)
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.LPALIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.LINKLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.DSSVM ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.SVCLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.NUCLEUS ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.DCMLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.INDMAC ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.CMDLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.HELP ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.SAMPLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.MACLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.PROCLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.TELCLIB ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.UADS ) VOL(SG2001) DEVT(3330))
  DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.VTAMLIB ) VOL(SG2001) DEVT(3330))

```

Figure 32 (Part 3 of 4). Sample Job Stream to Make the Master Catalog Recoverable

```

DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.IMAGELIB ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.PARMLIB ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.BROADCAST ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.MANX ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.MANY ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.DUMP00 ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.DUMP01 ) VOL(SG2001) DEVT(3330))
DEF CAT(AMASTCAT) NVSAM (NAME(SYS1.LOGREC ) VOL(SG2001) DEVT(3330))
/*
//STEP11A EXEC PGM=IEHPROGM,COND=(0,LT)
//SG2001 DD DISP=OLD,VOL=SER=SG2001,UNIT=3330
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
SCRATCH DSNAME=SYS1.NUCLEUS,VOL=3330=SG2001,MEMBER=SYSCATLG
/*
//STEP11B EXEC PGM=IEBGENER,COND=(0,LT)
//SYSIN DD DUMMY
//SYSPRINT DD SYSOUT=A
//SYSUT2 DD DSN=SYS1.NUCLEUS(SYSCATLG),DISP=(OLD,KEEP),
//DCB=BLKSIZE=13030,VOL=SER=SG2001,UNIT=3330
//SYSUT1 DD *
SG2001 (SEE NOTE 1)
/*
//
//JOB3 JOB ...
/* STEP 12 - REIPL THE SYSTEM
//STEP13 EXEC PGM=IDCAMS
//SYSPRINT DD SYSOUT=A
//CATPAK DD UNIT=3330,VOL=SER=CATPAK,DISP=OLD
//SYSIN DD *
ALTER AMASTCAT -
FILE(CATPAK) -
REMOVEVOLUMES(CATPAK)
/*
//

```

Notes:

1. This data record has the volume serial number of the volume containing the alternate catalog in columns 1 to 6.

Figure 32 (Part 4 of 4). Sample Job Stream to Make the Master Catalog Recoverable

GLOSSARY

The following terms are defined as they are used in this book. If you do not find the term you are looking for, refer to the index or to the Data Processing Glossary, GC20-1699.

Access Method Services: A multifunction service program that defines VSAM data sets 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.

addressed direct access: The retrieval or storage of a data record identified by its RBA, independent of the record's location relative to the previously retrieved or stored record. (See also keyed direct access, addressed sequential access, and keyed sequential access.)

addressed sequential address: The retrieval or storage of a data record in its entry sequence relative to the previously retrieved or stored record. (See also keyed sequential access, addressed direct access, and keyed direct access.)

alias: An alternative name for an entry.

alias entry: An entry that relates an alias (alternate entryname) to the real entryname of a user catalog or nonVSAM 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.)

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.

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. 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.)

bind: (verb) To keep a data set that has been staged from a mass-storage volume to a direct-access storage staging drive on the staging drive until the data set is closed.

catalog: (See master catalog and user catalog.)

catalog cleanup: A process that allows you to delete entries if their volume is no longer available; catalog cleanup also allows you to delete a catalog even though it is not empty. Catalog cleanup is a function of the DELETE command.

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).

catalog recovery area: (See CRA.)

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.

collating sequence: An ordering assigned to a set of items, such that any two sets in that assigned order can be collated. As used in this publication, the order defined by the System/370 8-bit code for alphabetic, numeric, and special characters.

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.

compression: (See key compression.)

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 won't fit in the original control interval.

CRA: Catalog recovery area. An entry-sequenced data set that exists on each volume owned by a recoverable catalog, including the volume on which the catalog resides. The CRA contains copies of the catalog's records, and can be used to recover a damaged catalog.

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

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 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.

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

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. (See also addressed direct access and keyed direct access.)

distributed free space: Space reserved within the control intervals of a key-sequenced data set for inserting new records into the data set in key sequence; also, whole control intervals reserved in a control area for the same purpose.

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 entryname 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.

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

exception exit: A user-written exit routine that attempts to recover from exceptional processing situations and is processed before the SYNAD exit routine gets control. The exception-exit routine is similar to the SYNAD error exit routine, except that the exception-exit routine can be tailored by the user for the data or index component of a VSAM data set or alternate index.

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.

external sort: Sorting of data records into a new sequence using a small amount of virtual storage and two

temporary sortfiles (entry-sequenced data sets) on a direct-access storage volume. (See also internal sort.)

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

free space: (See distributed free space.)

generation data group entry: An entry that permits nonVSAM data sets to be associated with other nonVSAM data sets as generation data sets.

generation data set: One of a collection of historically related nonVSAM 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.

generic name: A qualified name in which one qualifier is replaced by an asterisk; the generic name applies to all entries that match the qualifiers supplied in the generic name.

horizontal pointer: A pointer in an index record that gives the location of another index record in the same level that contains the next key in collating sequence; used for keyed sequential access.

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 replication: The use of an entire track of direct-access storage to contain as many copies of a single index record as possible; reduces rotational delay.

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

index upgrade: The process of updating an alternate index to reflect the changes made to the contents of its base cluster.

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. (See also external sort.)

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.

job catalog: A catalog made available for a job, by means of a JOBCAT DD statement, or for a job step, by means of a STEPCAT DD statement.

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 compression: The elimination of characters from the front and the back of a key that VSAM does not need to distinguish the key from the preceding or following key in an index record; reduces storage space for an index.

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 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.

keyed direct access: The retrieval or storage of a data record by use of an index that relates the record's key to its relative location in the data set, independent of the record's location relative to the previously retrieved or stored record. (See also addressed direct access, keyed sequential access, and addressed sequential access.)

keyed sequential access: The retrieval or storage of a data record in its key sequence relative to the previously retrieved or stored record, as defined by the sequence set of an index. (See also addressed sequential access, keyed direct access, and addressed direct access.)

mass sequential insertion: A technique VSAM uses for keyed sequential insertion of two or more records in sequence into a collating position in a data set: more efficient than inserting each record directly.

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.

nonunique keys: An alternate key value in a base-cluster's data record that equals the alternate key value of another data record. When an alternate index is created such that it allows nonunique keys, more than one data record might be located with the same key value.

nonVSAM entry: A catalog entry that describes a nonVSAM data set. A nonVSAM 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.

page space: A VS2 system data set. A page space is cataloged as an entry sequenced cluster (that is, the page space entry is similar to a cluster entry, and it points to a data entry).

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. (See also horizontal pointer and vertical pointer.)

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.)

prime index: The index component of a key-sequenced data set. (See also index and alternate index.)

prime key: (See key.)

qualified name: A name that is segmented by periods; each name segment is referred to as a qualifier.

random access: (See direct access.)

RBA: 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.

record: (See index record, data record, stored record.)

record replacement: An enhancement of the REPRO command which allows the user to merge two data sets and, when duplicate records are found (that is, the key values of two records are equal) the data record from the source data set replaces the data record in the target data set.

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.)

recovery volume: The first volume of a prime index, if the VSAM data set is key-sequenced cluster; otherwise, the first volume of the data set, if ESDS.

relative byte address: (See RBA.)

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.

replication: (See index replication.)

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

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 checking: The process of verifying the order of a set of records relative to some field's collating sequence.

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. (See also addressed sequential access and keyed sequential access.)

simple name: A qualifier of a qualified entryname or dsname. Simple names may be one to eight characters and, in a series, are separated from each other by a period.

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.

sort: (See external sort, internal sort, and collating sequence.)

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.

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

step catalog: (See job catalog.)

stored record: A data record, together with its control information, as stored in auxiliary storage.

upgrade set: All the alternate indexes that VSAM has been instructed to update whenever there is a change to the data component of the base cluster.

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.)

vertical pointer: A pointer in an index record of a given level that gives the location of an index record in the next lower level or the location of a control interval in the data set controlled by the index.

volume cleanup: The process of deleting all VSAM data spaces from a volume and removing a VSAM catalog's ownership of the volume, even though the volume contains VSAM data records.

volume entry: A catalog entry that describes a volume owned by the catalog. All VSAM data spaces on the volume are described in the volume entry and the catalog is no longer available.

INDEX

A

- abbreviations, parameter 5,381
 - Access Method Services commands 4
 - functional 4
 - modal 5
 - guide to 1
 - introduction to 4
 - invoking 11
 - access, read only 27,130,288
 - accessibility to secondary extents 33
 - adding records through the REPRO command 104
 - adding volumes 123
 - ADDVOLUMES parameter (ALTER) 123
 - ALIAS parameter 64
 - in DEFINE command 149
 - in DELETE command 272
 - in LISTCAT command 317
 - ALL parameter (LISTCAT) 318
 - example of output 422
 - ALL subparameter (EXPORTRA) 293
 - ALL subparameter (RESETCAT) 350
 - allocating space
 - by range of key values 195
 - for a catalog 51
 - on unmounted volumes 46
 - allocation
 - dynamic 9,46
 - of a catalog recovery area (CRA) 57
 - of a volume 16
 - of catalog's space 51
 - of user catalogs 21
 - ALLOCATION parameter (LISTCAT) 318
 - example of output 429
 - (see also BUFFERSPACE parameter)
 - ALTER command 71,119
 - allocation parameters 119
 - alternate index and path
 - attributes 119
 - catalog parameter 125
 - entry types to which each parameter applies 121
 - examples 136
 - format 119
 - generation-data-group attribute parameters 119
 - name parameters 119
 - parameter summary table 121
 - order of catalog use 72
 - Parameters 123
 - Optional parameters 123
 - Required parameters 123
 - protection and integrity parameters 119
 - removing VSAM ownership of a volume 16,36
 - volume cleanup 36,73
 - altering catalog entries 71
 - alternate index
 - backing up (copying) 104
 - building 67
 - defining 61,151
 - deleting 77,272
 - exporting 84
 - importing 84
 - listing its catalog entry 75
 - printing its contents 114
 - alternate name
 - creating 64,149
 - examples 151
 - deleting 77,272
 - example 282
 - alternate output data sets 11
 - ALTERNATEINDEX parameter 61
 - in DEFINE command 151
 - in DELETE command 272
 - in LISTCAT command 317
 - amendments, summary of vii
 - AN, print chain option 362
 - AREAS parameter (PARM) 362
 - ATTEMPTS parameter 30
 - in ALTER command 123
 - in DEFINE command
 - ALTERNATEINDEX 160
 - Catalog 254
 - CLUSTER 190
 - PAGESPACE 231
 - PATH 241
 - USERCATALOG 254
 - note for TSO users 254
 - attributes 58
 - altering 71
 - changing 71,482
 - defining 46,58
 - nullifying 131
 - AUTHORIZATION parameter
 - in ALTER command 124,129
 - in DEFINE command
 - ALTERNATEINDEX 160
 - CLUSTER 190
 - PAGESPACE 231
 - PATH 242
 - USERCATALOG 254
 - nullifying 129
 - Auxiliary Storage Management 235
 - authorization to process a data set
 - passwords 27
 - RACF protection 31
 - user-security-verification routine 31
- ### B
- backing up catalogs 35
 - backing up data sets 32,104
 - beginning location
 - in PRINT command 329
 - in REPRO command 338
 - BIND parameter 25
 - in ALTER command 133
 - in DEFINE command
 - used to define an alternate index 173
 - used to define a cluster 204
 - blanks vi,6
 - BLDINDEX command 67,138
 - catalog parameter 138
 - DD statements that describe the SORT workfiles 69
 - example 141
 - format 138
 - how an alternate index is built 67
 - order of catalog use 69
 - parameters 138

- Optional parameters 138
- Required parameters 138
- sort parameters 69
- bold face type, in notation conventions vi
- braces, in notation conventions vi
- brackets, in notation conventions vi
- BUFFERSPACE parameter
 - in ALTER command 124
 - in DEFINE command
 - used to define an alternate index 161
 - used to define a catalog 256
 - used to define a cluster 191
- building an alternate index 67,138

C

- calculating space for a catalog 51
- CALL macro, invoking Access Method Services with 14,475
- CANDIDATE parameter (DEFINE SPACE) 250
- capitalization, in notation conventions vi
- catalog
 - (see also master catalog and user catalog)
 - allocating user catalogs 51
 - altering 71,119
 - backing up 35,111
 - calculating space for 54
 - cleanup 38
 - connecting to master catalog (IMPORT) 82,299
 - converting devices 104
 - creating 253
 - defining 51
 - displaying 75
 - dumping a 35
 - duplicate names, preventing 46
 - entries, moving 82
 - identifying the volume of a 54
 - listing 75
 - master 20
 - modifying 71
 - order of use
 - ALTER 72
 - BLDINDEX 69
 - DEFINE 47
 - DELETE 80
 - LISTCAT 75
 - OS CVOLs (control volumes) 22
 - OS entries, converted to VSAM entries 23,86
 - ownership of a volume 16,36
 - preventing duplicate names 46
 - protecting 34
 - recovery 39,93
 - relationships among 21
 - reloading 112
 - resetting catalog entries 98
 - restoring the entries after system failure 35,93
 - secondary allocation amount 56
 - setting up the master catalog 20
 - space estimates 54
 - space, allocation of 51
 - transporting user catalogs 38
 - unloading 112
 - update password and nonVSAM data sets 54
 - use of, with data and space management 18
 - user catalog defined on a mass storage volume 51
 - using a backup copy 35
 - using a model to define 259
 - using qualified names for cataloged objects 18
 - VSAM's use of 18
- catalog entries in OS catalog converted 22,86
 - entry types and VSAM equivalents 87
- CATALOG parameter
 - in ALTER command 125
 - in BLDINDEX command 138
 - in CNVTCAT command 145
 - in DEFINE command
 - ALIAS 149
 - ALTERNATEINDEX 178
 - Catalog 263
 - CLUSTER 209
 - GENERATIONDATAGROUP 219
 - NONVSAM 226
 - PAGESPACE 237
 - PATH 247
 - SPACE 251
 - USERCATALOG 263
 - in DELETE command 273
 - in IMPORT command 301
 - in IMPORTRA command 310
 - in LISTCAT command 319
 - in LISTCRA command 324
 - in RESETCAT command 350
- catalog password protection 27
- catalog record
 - deleting 77
 - listing 75
 - modifying 71
 - resetting 98
 - restoring after system failure 93
- catalog recovery 39
- catalog recovery area 39,93
 - contents, description of 93
 - copying an entry from 95
 - listing the 95
 - examples of output 442
 - procedure for using 93
 - space allocation 57
- catalog use, order of, with
 - ALTER 72
 - BLDINDEX 69
 - DEFINE 47
 - DELETE 80
 - LISTCAT 75
- cataloging
 - aliases 64,149
 - alternate indexes 61
 - clusters 58
 - data sets 55,64
 - entry-sequenced data sets 55
 - general discussion of 45
 - key-sequenced data sets 55
 - nonVSAM data sets 64
 - user catalogs 18
- CHAIN parameter (PARM) 362
 - of master catalog 480
- changing attributes
 - allocation 119
 - generation data group 119
 - name 119
 - protection 119
 - VSAM master catalog 480

changing a volume's serial number 37
 changing the attributes of a VSAM
 master catalog 480
 CHARACTER parameter (PRINT) 328
 CHECKID parameter (CHKLIST) 142
 checkpoint data set 116
 checkpoint, tape data sets open
 at 116
 CHKLIST command 116,142
 examples 143
 format 142
 listing tape volumes mounted at
 checkpoint 116
 parameters 142
 optional parameters 142
 required parameters 142
 sample output from 455
 CHKPT JCL parameter 438
 cleanup
 catalog 38,270
 volume 36
 cluster
 backing up (copying) 32,104
 defining 58,182
 deleting 77,270
 exporting 82,286
 extending volumes xxxx
 listing its entry 75,316
 importing 82,299
 printing its contents 114,327
 verifying its end-of-file
 values 91,356
 CLUSTER parameter
 in DEFINE command 182
 in DELETE command 272
 in LISTCAT command 317
 cluster type, specifying 184
 CNVTCAT command 86,145
 examples 146
 format 145
 parameters 145
 CODE parameter 29
 in ALTER command 125
 in DEFINE command
 ALTERNATEINDEX 161
 CATALOG 256
 CLUSTER 191
 PAGESPACE 232
 PATH 243
 USERCATALOG 256
 nullifying 129
 command continuation 7
 command execution, controlling 358
 command reference summary 455
 command statement syntax 5
 commands
 ALTER 71,119
 BLDINDEX 67,138
 CHKLIST 116,142
 CNVTCAT 86,145
 continuation of 7
 DEFINE
 ALIAS 64,149
 ALTERNATEINDEX 61,151
 CLUSTER 58,182
 for an entry-sequenced
 cluster 189
 for a key-sequenced
 cluster 186,189
 for a relative-record
 cluster 189
 GENERATIONDATAGROUP 65,218
 NONVSAM 64,223
 PAGESPACE 65,228
 PATH 63,240
 SPACE 57,249
 USERCATALOG 51,253
 DELETE 77,270
 DO 5, 360
 END 5,360
 EXPORT 82,286
 EXPORTRA 93,293
 functional 4
 general structure 5
 guide to 1
 IF 5,359
 IMPORT 82,299
 IMPORTRA 93,309
 introduction 4
 LISTCAT 75,316
 LISTCRA 95,324
 modal 5,358
 not allowed under TSO 13
 parameter set 5
 reference summary 455
 PARM 5,361
 PRINT 114,327
 REPRO 104,337
 RESETCAT 98,350
 SET 361
 VERIFY 91,356
 structure of 5
 terminator 8
 types 4
 commas 6
 comments 6
 COMPARE parameter (LISTCRA) 325
 example of output 453
 components parameters (DEFINE)
 used to define an alternate
 index 176
 used to define a catalog 262
 used to define a cluster 206
 concatenated DD statements 9
 condition codes 358
 conditional command execution
 using LASTCC 361
 using MAXCC 361
 conditional statements 358
 CONNECT parameter (IMPORT) 299
 connecting a user catalog to the
 master catalog 20,64,82
 considerations, language 5
 continuation cautions 7
 continuing commands 7
 for multivolume data sets
 (RESETCAT) 101
 job control language 9
 parameter set 5
 terminator 8
 continuation of commands 7
 control-access password 27
 (see also CONTROLPW parameter)
 control area
 in catalogs 51
 preformatting 32
 control volume pointer entries,
 converted 22,86
 control volumes, OS 20,86
 converting to VSAM catalogs 86
 relationship to other catalogs 20
 CONTROLINTERVALSIZE parameter (DEFINE)
 used to define an alternate
 index 162
 used to define a cluster 192
 controlling command execution 358
 CONTROLPW parameter
 in ALTER command 125

- in ALTER command (NULLIFY) 129
- in DEFINE command
 - ALTERNATEINDEX 162
 - CLUSTER 193
 - PAGESPACE 232
 - PATH 243
 - USERCATALOG 258
- nullifying 129
- conventions
 - notation vi
 - syntactical 5
- converting catalog entries 86
- converting data sets 104
- converting devices for a user
 - catalog 104
- copying
 - alternateindexes 29,104
 - catalogs 110
 - example 344
 - for backup 35,111
 - for device conversion 104
 - example 344
 - data sets 1,104
 - from the catalog recovery area 95
 - copying and printing, example of 375
 - correcting end-of-data-set
 - information 91
 - correcting end-of-key-range
 - information 91
 - corrective measures
 - for catalog integrity 40
 - for data integrity 34
- COUNT parameter
 - in PRINT command 331
 - in REPRO command 341
- CRA (see catalog recovery area)
- CRA parameter (EXPORTRA) 293
- CRAFILES parameter (RESETCAT) 350
- CRAVOLUMES parameter (RESETCAT) 350
- creating
 - alias names 64,149
 - alternate indexes 67,138
 - an alternate index and path, example
 - of 383
 - catalogs 51,253
 - clusters 58,153,156
 - data sets 58,64,182
 - data spaces 57,249
 - entry-sequenced data sets 58
 - generation data groups 65,218
 - key-sequenced data sets 58
 - master catalog 51,253
 - nonVSAM data sets 64
 - page spaces 65,228
 - path 63,240
 - portable data sets 82
 - user catalogs 51,253
- CREATION parameter (LISTCAT) 319
 - example of output 434
- cross-region sharing 200
- cross-system sharing 200
- CVOLEQUATES parameter (CNVTCAT) 145
- CVOLs (control volumes), OS 22,86
 - converting to VSAM catalogs 86
 - relationship to other catalogs 21
- CYLINDERFAULT parameter 25
 - in ALTER command 134
 - in DEFINE command
 - used to define an alternate
 - index 173
 - used to define a cluster 204
- CYLINDERS parameter
 - used to define an alternate
 - index 157

- used to define a cluster 187
- used to define a data space 250
- used to define a page space 229
- used to define a user catalog 253

D

- data, storage of 15
- data integrity
 - backing up data 32
- data organization, specifying 189
- DATA parameter
 - in DEFINE command
 - ALTERNATEINDEX 176
 - CATALOG 262
 - CLUSTER 206
 - USERCATALOG 263
 - in LISTCAT command 317
- data portability 21,33,82
- data protection 32
- data security
 - and protection 27
 - authorization routine 14,27,31
 - passwords 27
- data set
 - allocation 46
 - backup copy 32,82,104
 - catalog entry 45
 - cataloging 45,58,64
 - copying 33,104
 - comparison between types of 87
 - creating 40,58,64
 - defining 40,58,64
 - deleting 77
 - entry-sequenced 58
 - generated names 17
 - generation 218
 - key-sequenced 58
 - listing 114
 - loading records into 111
 - NONVSAM 64
 - organization 17,189
 - output 11
 - portability 33,82
 - preformatting 32
 - reorganizing 112
 - tape, open at checkpoint 116
 - transporting 32,82
 - unique 51,205,212
 - using qualified names that identify
 - the catalog 20,22
- data-set access 288
- data-set security bits in Format-1
 - DSCB 17
- data-set type, specifying 64,189
- data space
 - allocating the volume for 57,249
 - defining 249
 - examples 251
 - extending 249
 - pagespace in 228
 - VSAM objects in a 16,58
- data sets, VSAM 16,58
- data spaces on a volume 16
- data structure, VSAM 15
- DATASET parameter (VERIFY) 356
- DD statements 6(see JCL)
 - for Access Method Services jobs 10
 - for JOBCAT and STEPCAT
 - catalogs 21,46
 - for output data sets 11

- for RESETCAT work file 101
- for the BLDINDEX sort workfiles 69
- debugging tool 361
- default margins 5
- default output data set 11
- define a user catalog, example of 365
- define nonVSAM data sets, example of 371
- define VSAM data sets, example of 367
- DEFINE command 45
 - how parameters are organized 48
 - job control language statements for 46
 - naming entries to prevent duplicate names 46
 - order of catalog use 47
 - used to define an alternate index 61,151
 - catalog parameter 178
 - components parameters 176
 - data integrity parameters 155
 - example 179
 - format 151
 - index parameters 169
 - model parameter 167
 - name parameter 154,156
 - optional parameters 160
 - parameters, summary of 154
 - required parameters 156
 - protection parameters 155
 - specifying information for 61,151
 - used to define an alternate name (alias) 64,149
 - example 149
 - format 149
 - parameters 149
 - optional parameters 149
 - required parameters 149
 - used to define a catalog
 - allocation parameters 240
 - catalog parameter 247
 - example 247
 - format 240
 - model parameter 244
 - name parameter 241
 - parameter summary 240
 - optional parameters 241
 - required parameters 241
 - protection parameters 240
 - update parameter 241
 - used to define a user
 - catalog 51,253
 - allocation parameters 253
 - catalog parameter 253,263
 - components parameters 253,262
 - entry type parameters 253
 - examples 264
 - format 253
 - model parameter 253,259
 - name parameter 253
 - parameters, summary of 253
 - optional parameters 254
 - required parameters 253
 - protection and integrity parameters 253
 - used to define a cluster 58,182
 - allocation parameters 184
 - catalog parameter 209
 - components parameters 206
 - data organization parameters 189
 - entry type parameter 184
 - examples 209
 - format 182
 - model parameter 197
 - name parameter 184
 - parameters, summary of 184
 - optional parameters 190
 - required parameters 186
 - parameters that apply only to key-sequenced clusters 186
 - protection and integrity parameters 185
 - specifying information for 58,182
 - used to define a data space 57,249
 - allocating the volume 57
 - allocation parameters 249
 - catalog parameter 251
 - example 252
 - format 249
 - optional parameters 251
 - required parameters 249
 - used to define a generation data group 65,218
 - example 220
 - format 218
 - parameters 218
 - used to define a nonVSAM data set 64,223
 - example 227
 - format 223
 - parameters 223
 - optional parameters 224
 - required parameters 223
 - used to define a page space 65,228
 - allocation parameters 197
 - catalog parameter 237
 - examples 238
 - format 228
 - model parameter 234
 - name parameter 228
 - parameter summary 228
 - optional parameters 231
 - required parameters 229
 - parameters that apply only to page spaces 229
 - protection and integrity parameters 229
 - used to define a path 63,240
 - allocation parameter 240
 - catalog parameter 247
 - example 247
 - format 240
 - model parameter 244
 - name parameter 241
 - parameter summary 240
 - optional parameters 241
 - required parameters 241
 - protection parameters 240
 - update parameter 241
 - used to define a user
 - catalog 51,253
 - allocation parameters 253
 - catalog parameter 253,263
 - components parameters 253,262
 - entry type parameters 253
 - examples 264
 - format 253
 - model parameter 253,259
 - name parameter 253
 - parameters, summary of 253
 - optional parameters 254
 - required parameters 253
 - protection and integrity parameters 253
 - How parameters are organized 48
 - defining 45
 - alias 64,149
 - alternate index 61,151

- catalog 51,253
- cluster 58,182
- data space 57,249
- generation data group 65,218
- NONVSAM data set 64,223
- page space 65,228
- path 63,240
- user catalog 51,253
- DELETE command 270
 - catalog cleanup 38,278
 - examples 277
 - format 270
 - kinds of catalog entries you can delete 77
 - parameters 270
 - optional parameters 272
 - required parameters 270
 - removing VSAM ownership of a volume 15
- deleting 77
 - (see also DELETE command)
 - catalog entries 77
 - a data set 78
 - examples of 277
 - a VSAM catalog and its objects, example of 391
 - regardless of retention date (see PURGE parameter)
- descriptive information, specifying for an alternate index 61
 - for a cluster 58
- DESTAGEWAIT parameter 25
 - in ALTER command 126
 - in DEFINE command
 - ALTERNATEINDEX 163
 - catalog 258
 - CLUSTER 193
 - USERCATALOG 258
- destaging attributes 25
 - (see also DESTAGEWAIT and NODESTAGEWAIT)
- destaging to mass storage 25
- device conversion for a user catalog 104
- device independence of VSAM data sets 16
- device type, esoteric 223
- device type translate table 414
- DEVICETYPE subparameter
 - in IMPORT command 303
 - in IMPORTRA command 311
- DEVICETYPES parameter (DEFINE NONVSAM) 223
- direct-access space allocation 46
- DISCONNECT parameter (EXPORT) 286
- displaying catalog information (see LISTCAT command)
- DO, in DO-END command sequence 361
 - examples 364
 - format 361
- DSCB (data set control block)
 - volume cleanup 36
 - VSAM volume ownership 16
- DUMMY subparameter (REPRO) 338
- dummy records 338
- DUMP parameter
 - in LISTCRA command 326
 - example of output 450
 - in PRINT command 328
- dump points 363
- DUMP/RESTORE program, IEHDASDR
 - backing up a catalog 35
 - timestamp updating 16
- dumping

- a catalog and its data sets 35
 - a volume owned by VSAM 35
- duplicate names in a catalog, preventing 46
- dynamic allocation 9,46

E

- ellipses, in notation conventions vi
- ELSE clause, in IF-THEN-ELSE command sequence 361
- EMPTY parameter
 - in ALTER command 126
 - in DEFINE command
 - used to define a generation data group 218
- END, in DO-END command sequence 361
 - examples 364
 - format 361
- end-of-file values, restoring a cluster's 91
- ending location
 - in PRINT command 331
 - in REPRO command 341
- enhancements, summary of vii
- ENTRIES parameter (LISTCAT) 320
- ENTRIES subparameter (EXPORTRA) 293
- entryname/password parameter
 - in ALTER command 123
 - in DELETE command 270
 - in EXPORT command 286
- entry-sequenced data set 15,58
 - compared with key-sequenced 15,58
 - examples of defining 210
 - specification of 59
- entry-type parameters (LISTCAT) 318
- entry-types each ALTER parameter applies to 121
- ENVIRONMENT subparameter (REPRO) 338
- EQ (equal) 360
- ERASE parameter
 - erasing data on a staging drive 25
 - in ALTER command 126
 - in DEFINE command
 - used to define an alternate index 163
 - used to define a cluster 193
 - in DELETE command 274
 - in EXPORT command 287
 - in IMPORT command 302
 - erasing (deleting) a data set 77
 - (see also DELETE command and ERASE parameter)
- estimating the space requirements
 - for a catalog 54
 - example of 264
 - worksheet for 55
 - for the BLDINDEX sort workfiles 68
- examples of
 - Access Method Services
 - jobstream 366
 - ALTER command 136
 - BLDINDEX command 141
 - catalog unload/reload 348
 - CHKLIST command 143,455
 - CNVTCAT command 146
 - copying a catalog 345
 - DEFINE command
 - for an alias 151
 - for an alternate index 179
 - for a catalog 264

- for a cluster 210
- for a data space 252
- for a generation data group 220
- for a nonVSAM data set 227
- for a page space 238
- for a path 247
- for a user catalog 264
- DELETE command 277
- DO-END 364
- EXPORT command 290
- EXPORTRA command 296
- IF-THEN-ELSE 364
- IMPORT command 308
- IMPORTRA command 315
- jobs using Access Method Services commands 366
- LISTCAT command 323
- LISTCAT output 414,436
- LISTCRA command 327
- LISTCRA output 441
- PRINT command 333
- PRINT output 335
- REPRO command 343
- RESETCAT command 353
- SET 364
- VERIFY command 358
- EXCEPTIONEXIT parameter
 - in ALTER command 126
 - in ALTER (NULLIFY) 129
 - in DEFINE command
 - used to define an alternate index 164
 - used to define a cluster 194
- execution, controlling 359
- expiration date, nullifying 130
 - (see also FOR parameter and TO parameter)
- EXPIRATION parameter (LISTCAT) 322
 - example of output 434
- EXPORT command 286
 - examples 290
 - format 286
 - parameters 286
 - optional parameters 287
 - required parameters 286
 - transporting alternate indexes 33
 - transporting data sets 33
 - transporting user catalogs 21
- exporting
 - a base cluster and its alternate index, example of 386
 - a catalog's data sets 33
 - an entry 82
 - a user catalog 21
 - examples 290
- EXPORTRA command 293
 - copying an entry from the catalog recovery area 95
 - example of 296
 - format 295
 - parameters 251
 - optional parameters 295
 - required parameters 293
- EXTERNALSORT parameter (BLDINDEX) 139

F

- field continuation rules 7
- fields
 - continuing 7
 - separators with 7

- FILE parameter
 - (see also INFILE parameter and OUTFILE parameter)
 - in ALTER command 126
 - in DEFINE command
 - to define an alternate index 164
 - to define a cluster 194
 - to define a data space 251
 - to define a page space 233
 - to define a path 244
 - to define a user catalog 258
 - in DELETE command 274
 - in VERIFY command 357
- FILE subparameter
 - in IMPORT command 303
 - in IMPORTRA command 312
- FILESEQUENCENUMBERS parameter (DEFINE NONVSAM) 224
- FOR parameter
 - in ALTER command 134
 - in DEFINE command
 - ALTERNATEINDEX 174
 - catalog 261
 - CLUSTER 205
 - GENERATIONDATAGROUP 218
 - NONVSAM 226
 - PAGESPACE 235
 - PATH 246
 - USERCATALOG 261
- FORCE parameter
 - in DELETE command 275
 - in EXPORTRA command 295
- format of commands
 - ALTER command 119
 - BLDINDEX command 138
 - CHKLIST command 142
 - CNVTCAT command 145
 - DEFINE command
 - ALIAS 149
 - ALTERNATEINDEX 151
 - CLUSTER 182
 - GENERATIONDATAGROUP 218
 - NONVSAM 223
 - PAGESPACE 228
 - PATH 240
 - SPACE 249
 - USERCATALOG 253
 - DELETE command 270
 - DO 361
 - END 361
 - EXPORT command 286
 - EXPORTRA command 293
 - IF 360
 - IMPORT command 300
 - IMPORTRA command 310
 - LISTCAT command 317
 - LISTCRA command 325
 - PARM 362
 - PRINT command 328
 - REPRO command 338
 - RESETCAT command 351
 - SET 362
 - VERIFY command 357
- format of LISTCAT output 414
- format of LISTCRA output 438
- format of printed data with PRINT command 329
- Format 1 DSCB 16,36
- Format 4 DSCB 16,36
- FREESPACE parameter
 - in ALTER command 127
 - in DEFINE command
 - used to define an alternate index 165

- used to define a cluster 194
- FROMADDRESS parameter
 - in PRINT command 330
 - in REPRO command 340
- FROMKEY parameter
 - in PRINT command 330
 - in REPRO command 339
- FROMNUMBER parameter
 - in PRINT command 330
 - in REPRO command 340
- FULL parameter (PARM) 363
- functional commands 4

G

- GE (greater than or equal) 360
- generated names 17
 - for suballocated data spaces 17
 - for unique data spaces 17
 - for VSAM data and index components 17
- generating trace tables 362
- generation data set 218
 - cataloging, example of 220
- GENERATIONDATAGROUP parameter
 - in DEFINE command 218
 - example 220
 - in DELETE command 272
 - in LISTCAT command 318
- generic (qualified) names 22
 - in ALTER command 72
 - in DEFINE command 46
 - example 216
 - in DELETE command 81
 - in LISTCAT command 75
- GRAPHICS parameter 363
- GT (greater than) 360
- guide to Access Method Services 1

H

- HEX parameter (PRINT) 329
- HISTORY parameter (LISTCAT) 319
 - example of output 432
- HN, print chain option 363
- how an alternate index is built 67
- how to use one object as a model for another object 49
- hyphens 7

I

- I/O buffer
 - and the ALTER command 124
 - and the ALTERNATEINDEX command 161
 - and the DEFINE command 191,256
 - space for 161,191,256
- IDCAMS 4,12
- identifying the catalog's volume 54
- IEHDASDR DUMP/RESTORE program
 - backing up a catalog 35
 - changing a volume's serial number 37
- IF-THEN-ELSE command sequence 360
 - examples 363

- format 360
- IGNORE parameter (RESETCAT) 351
- IMBED parameter (DEFINE)
 - used to define an alternate index 165
 - used to define a cluster 195
- INTOEMPTY parameter
 - in IMPORT command 302
- IMPORT command 33,82,300
 - examples 308
 - format 300
 - parameters 300
 - optional parameters 302
 - required parameters 300
 - transporting alternate indexes 33
 - transporting data sets 33
 - transporting user catalogs 21
- importing
 - a catalog's data sets 33
 - an entry 82
 - a user catalog 21
 - examples 308
- IMPORTRA command 310
 - example of 315
 - format 310
 - parameters 310
 - restoring the catalog entry that was obtained using the EXPORTRA command 97
- INDATASET parameter
 - in CNVTCAT command 145
 - in IMPORT command 301
 - in PRINT command 328
 - in REPRO command 338
- index, alternate (see alternate index)
- INDEX parameter
 - in DEFINE command
 - used to define an alternate index 177
 - used to define a catalog 263
 - used to define a cluster 208
 - used to define a user catalog 263
 - in LISTCAT command 318
- INDEXED parameter (DEFINE) 189
- INFILE parameter
 - in BLDINDEX command 138
 - in CHKLIST command 142
 - in CNVTCAT command 145
 - in EXPORT command 287
 - in IMPORT command 300
 - in IMPORTRA command 310
 - in LISTCRA command 325
 - in PRINT command 328
 - in REPRO command 338
- INFILE subparameter (EXPORTRA) 294
- INHIBIT parameter (ALTER) 127
- INHIBITSOURCE parameter (EXPORT) 288
- INHIBITTARGET parameter (EXPORT) 288
- input record margins 6
- integrity of data
 - passwords 27
- INTERNALSORT parameter (BLDINDEX) 139
- interpreting LISTCAT output
 - listings 395
 - description of keyword fields 402
 - LISTCAT output keywords 395
 - LISTCAT output examples 414
- interpreting LISTCRA output 441
- introduction 4
- invoking Access Method Services 11,472
 - as a job or job step 13
 - from a processing program 15,472
 - from a TSO terminal 8

invocation from a PL/I program 477
ISAM data sets
 cataloging 64
 converting 104
 copying 104
 deleting 77
 listing catalog entries for 75
 printing 114
italics, in notation conventions vi

J

JCL
 allocating user catalogs 21
 catalog unload/reload 114
 CHKLIST command 143
 DEFINE command 46
 dynamic allocation 9,46
 for an output data set 11
 example 414
 identifying VSAM data sets or
 volumes 17
 invoking Access Method Services 11
 examples of use 366
 JOBSTAT DD statement 21,110
 LISTCAT command 414
 passwords for nonVSAM data sets 30
 RESETCAT command 102
 STEPSTAT DD statement 21,110
 to be avoided 438
 to define an entry 46
JOBSTAT catalog 110
 DD statement for 21,110
 order of catalog use
 ALTER 72
 BLDINDEX 69
 DEFINE 47
 DELETE 80
 LISTCAT 75
jobstep, invoking Access Method
 Services as a 11

K

key
 allocating space on volumes by key
 range 195
 (see also KEYS parameter)
key-pointer pairs of an alternate
 index 67
 building 67
 defining 67
key-sequenced data set
 definition 16
 examples of defining 210,212,215
 specification of 189
KEYRANGES parameter (DEFINE)
 used to define an alternate
 index 165
 used to define a cluster 195
KEYRANGES subparameter (IMPORT) 165
KEYS parameter
 in ALTER command 127
 in DEFINE command
 used to define an alternate
 index 167
 used to define a cluster 196
keyword parameters 5

kinds of catalog entries you can
 delete 77

L

language considerations 5
 continuation cautions 7
 continuing commands 7
 parameter set 5
 terminator 8
last condition code (see LASTCC)
LASTCC
 in IF 360
 in SET 362
LE (less than or equal) 360
levels in LISTCAT command 319
LEVEL parameter (LISTCAT) 319
LIMIT parameter (DEFINE
 GENERATIONDATAGROUP) 218
LIST parameter (CNVTCAT) 145
LISTCAT command 75,317
 examples 323
 format 317
 interpreting output from 395
 JCL required for 414
 order of catalog use 75
 parameters 317
 optional parameters 317
 required parameters 317
 sample output from 414
LISTCAT output
 and Access Method Services
 messages 416
 keywords
 description of 402
 list of 395
 job control language for 414
 listing, examples of
 LISTCAT output 414
 LISTCAT ALL output 422
 LISTCAT ALLOCATION output 429
 LISTCAT CREATION/EXPIRATION
 output 434
 LISTCAT HISTORY output 432
 LISTCAT SPACE ALL output 421
 LISTCAT VOLUME output 419
LISTCRA command 325
 example 327
 examples of output 441
 format 325
 listing the catalog recovery area's
 contents 93
 parameters 325
 optional parameters 325
 required parameters 325
 types of listing 442
LISTCRA output 441
 description of 441
 listing, examples of
 LISTCRA DUMP COMPARE output 448
 LISTCRA DUMP NOCOMPARE output 450
 LISTCRA NAME COMPARE output 452
 LISTCRA NAME NOCOMPARE output 448
listing catalog entries 75
 examples of 323
 examples of output 414
listing tape volumes mounted at
 checkpoint 116
listing the catalog recovery area's
 contents 95
 examples of output 441

- loading catalogs (unload/reload) 112
- loading records into a data set 106
 - preformatting control areas 32
 - REPRO command 338
- loading VSAM data sets 106
- lower case, in notation
 - conventions vi
- LRECL 104
- LT (less than) 360

M

- making a copy of a data set 33
- making a user catalog available 21
- MARGINS parameter (PARM) 363
- margins, default 6
- Mass Storage System, IBM 3860 25
- mass storage volume 25
- master catalog 20
 - (see also user catalog)
 - backing up 111
 - cataloging nonVSAM data sets 64
 - changing the attributes of 535
 - creating a 51
 - deleting restriction 78
 - improving availability of 35
 - order of catalog search
 - ALTER 72
 - BLDINDEX 69
 - DEFINE 47
 - DELETE 80
 - LISTCAT 75
 - preventing duplicate names in 46
 - protecting 34
 - reloading 112
 - relationship to user catalogs 19,51
 - unloading 112
 - using a backup copy 35
- master password 27
 - (see also MASTERPW parameter)
- MASTERCATALOG parameter
 - in CNVICAT command 145
 - in DEFINE command (defines user catalog) 51
- MASTERPW parameter
 - in ALTER command 128
 - in DEFINE command
 - ALTERNATEINDEX 167
 - CLUSTER 197
 - PAGESPACE 233
 - PATH 244
 - USERCATALOG 259
 - in EXPORTRA command 295
 - in LISTCRA command 325
 - in RESETCAT command 351
 - nullifying 130
- MAXCC parameter
 - in IF 360
 - in SET 362
- maximum condition code 360,362
- memory (see virtual storage)
- merging VSAM data sets 104
- migration of passwords 28
- MISCOMPARES messages 42,441
- mismatches between catalog and CRA records 40
- modal commands 5
 - not allowed from TSO terminal 14
- MODEL parameter
 - to define an alternate index 167
 - to define a catalog 259

- to define a cluster 197
- to define a page space 234
- to define a path 244
- to define a user catalog 259
- modeling 49
 - alternate indexes 167
 - clusters 197
 - example 216
 - data components 49
 - index components 49
 - page spaces 234
 - paths 244
 - user catalogs 259
 - example 267
- modify sequence of execution 360
- modifying catalog information (see ALTER command)
- MODULE subparameter (ALTER) 129
- moving data sets between systems 35,82
- moving entries 82
- moving user catalogs between systems 21,83

N

- name parameter (ALTER) (see entryname parameter)
- NAME parameter
 - in DEFINE command
 - used to catalog a nonVSAM data set 223
 - used to identify an alternate index 156
 - used to identify a catalog 254
 - used to identify a cluster 186
 - used to identify a generation data group 218
 - used to identify a page space 229
 - used to identify a path 241
 - in LISTCAT command 319
 - example of output 414
 - in LISTCRA command 325
 - example of output 448,452
- names, generated 16
- names, generic (or qualified)
 - in ALTER command 72
 - in DEFINE command 46
 - example 216
 - in DELETE command 81
 - in LISTCAT command 75
- naming
 - clusters 186
 - data components 186
 - entry-sequenced data sets 186
 - index components 186
 - key-sequenced data sets 186
 - master catalog 46
 - page spaces 229
 - paths 241
 - user catalogs 46,218
- NE (not equal) 360
- NEWNAME parameter (ALTER) 128
- NEWNAME subparameter (IMPORT) 305
- NOCOMPARE parameter (LISTCRA) 325
 - example of output 448,450
- NODESTAGEWAIT parameter 25
 - in ALTER command 126
 - in DEFINE command
 - ALTERNATEINDEX 163
 - catalog 258

CLUSTER 193
 USERCATALOG 258
 NOEMPTY parameter
 in ALTER command 126
 in DEFINE command
 used to define a generation data group 218
 NOERASE parameter
 in ALTER command 126
 in DEFINE command
 used to define an alternate index 163
 used to define a cluster 193
 in DELETE command 274
 in EXPORT command 287
 in IMPORT command 302
 not erasing data on a staging drive 25
 NOFORCE parameter
 in DELETE command 275
 in EXPORTRA command 295
 NOIGNORE parameter (RESETCAT) 351
 NOIMBED parameter (DEFINE)
 used to define an alternate index 165
 used to define a cluster 195
 NOINHIBITSOURCE parameter (EXPORT) 288
 NOINHIBITTARGET parameter (EXPORT) 288
 NOLIST parameter (CNVTCAT) 145
 NONE subparameter (EXPORTRA) 295
 NONE subparameter (RESETCAT) 350
 NONINDEXED parameter (DEFINE) 190
 NONSPANNED parameter (DEFINE) 203
 NONUNIQUEKEY parameter
 in ALTER command 129
 in DEFINE command 174
 nonVSAM data sets
 and the catalog's update password 30,54
 cataloging 64
 example 227
 copying 33
 deleting 77
 example of defining 227
 listing catalog entries for 75
 passwords for 30
 printing 114,328
 NONVSAM parameter
 in DEFINE command 64,223
 in DELETE command 272
 in LISTCAT command 318
 NOPURGE parameter
 in DELETE command 276
 in EXPORT command 288
 in IMPORT command 307
 NORECLAIM parameter
 in DEFINE command
 used to define an alternate index 169
 used to define a cluster 201
 NOREPLACE parameter (REPRO) 337
 NOREPLICATE parameter (DEFINE)
 used to define an alternate index 170
 used to define a cluster 201
 NOREUSE parameter
 in DEFINE command
 used to define an alternate index 170
 used to define a cluster 201
 in REPRO command 341
 normal output data set 11

NOSAVRAC parameter
 in IMPORT command 307
 in IMPORTRA command 314
 NOSCRATCH parameter
 in DEFINE command
 used to define a generation data group 218
 in DELETE command 276
 catalog cleanup 30,276
 modified in ALTER command 132
 notation conventions vi
 NOSWAP parameter (DEFINE PAGESPACE) 229
 in DEFINE PAGESPACE command 235
 in DEFINE PAGESPACE example 238
 NOTUSABLE parameter (LISTCAT) 322
 NOUPDATE parameter
 in ALTER command 134
 in DEFINE command 246
 NOUPGRADE parameter
 in ALTER command 135
 in DEFINE command 176
 NOTRECOVERABLE parameter (DEFINE) 260
 NOWRITECHECK parameter
 in ALTER command 135
 in DEFINE command
 ALTERNATEINDEX 176
 CLUSTER 206
 USERCATALOG 262
 NULLIFY parameter (ALTER) 129
 NUMBERED parameter (DEFINE) 190



objects (VSAM)
 in a data space 16,58
 space assignment to 58
 OBJECTS parameter
 in IMPORT command 303
 in IMPORTRA command 311
 OFF parameter (PARM) 363
 operator entering passwords 30
 (see also CODE parameter and ATTEMPTS parameter)
 optimizing the performance of catalog
 unload/reload 114
 optional allocation parameters
 for an alternate index 155
 for a cluster 185
 OR sign (|), in notation conventions vi
 order of catalog use
 ALTER 72
 BLDINDEX 69
 DEFINE 47
 DELETE 80
 LISTCAT 75
 ORDERED parameter (DEFINE)
 used to define an alternate index 167
 used to define a cluster 198
 ORDERED subparameter (IMPORT) 305
 OS CVOLs (control volumes)
 converting to VSAM catalogs 86
 relationship to other catalogs 19
 OUTDATASET parameter
 in EXPORT command 287
 in IMPORT command 302
 in REPRO command 339
 OUTFILE parameter
 in BLDINDEX command 138

- in CHKLIST command 142
- in EXPORT command 287
- in EXPORTRA command 292
- in IMPORT command 301
- in IMPORTRA command 314
- in LISTCAT command 322
- in LISTCRA command 325
- in PRINT command 331
- in REPRO command 339
- output data sets
 - default for listing 12
 - job control language for 12,444
 - requirements 11
- overwriting (see ERASE parameter)
- OWNER parameter
 - in ALTER command 130
 - in ALTER (NULLIFY) 130
 - in DEFINE command
 - note for TSO users 198
 - used to define an alternate index 168
 - used to define a catalog 260
 - used to define a cluster 198
 - used to define a generation data group 218
 - used to define a nonVSAM data set 227
 - used to define a page space 234
 - used to define a path 245
- ownership bit in Format-4 DSCB 17
- ownership, VSAM volume 17
 - removing 17

P

- PAGESPACE parameter
 - in DEFINE command
 - used to define a page space 57,228
 - examples 238
 - in DELETE command 272
 - in LISTCAT command 318
- parameter abbreviations 5
- parameter organization (DEFINE command) 48
- parameter set 5
- parameters
 - entry-types each ALTER parameter applies to 121
 - keyword 5
 - lists 6
 - parentheses within 6
 - positional 5
 - separators and 6
 - summary of
 - for the definition of an alternate index 154
 - for the definition of a catalog 253
 - for the definition of a cluster 184
 - for the definition of a user catalog 253
 - that apply only to key-sequenced clusters 186
- parentheses, in notational conventions vi
- PARM command 362
- password
 - control 27
 - (see also CONTROLPW parameter)

- for nonVSAM data set 30
- given by operator 30
- levels of authorization 27
- master 27
 - (see also MASTERPW parameter)
- read 27
 - (see also READPW parameter)
- update 27
 - (see also UPDATEPW parameter)
- path
 - defining 63,240
 - deleting 77
 - listing its catalog entry 75
- PATH parameter
 - in DEFINE command 63,240
 - in DELETE command 273
 - in LISTCAT command 318
- PATHENTRY parameter (DEFINE) 241
- performance, optimizing VSAM's catalog unload/reload 114
 - staging/destaging with mass storage 25
- performance options information, specifying
 - for an alternate index 62
 - for a cluster 60
- PERMANENT parameter (EXPORT) 289
- plus sign 6
- PN, print chain option 363
- portable data sets 33,82
- position, key 5
 - (see also KEYS parameter)
- positional parameters 5
- preformatting control areas 32
- preventing duplicate names in a catalog 46
- preventive measures for protecting catalogs 39
- primary allocation (see CYLINDERS parameter, RECORDS parameter, and TRACKS parameter)
- print chain options 363
- PRINT command 114,328
 - examples 333
 - format 328
 - parameters 328
 - optional parameters 329
 - required parameters 328
 - sample output from 329
- print graphics 363
- printing
 - catalog entries 75
 - catalog recovery area contents 95
 - data sets 114
 - names of tape data sets open at checkpoint 116
- printing and copying, an example of 376
- processing program, invoking Access Method Services from 12
- prompting codes for operator entering passwords 30
 - (see also CODE parameter and ATTEMPTS parameter)
- protecting
 - (see also data integrity and data security)
 - alternate indexes 62
 - catalogs 34
 - clusters 32,60
 - data 32
 - data components 32
 - index components 32
 - paths 240

- user catalogs 32
- protection and integrity parameters, specifying
 - for an alternate index 62
 - for a cluster 60
- protection parameters 60
 - in ALTER command 119
 - in DEFINE command
 - used to define an alternate index 155
 - used to define a catalog 253
 - used to define a cluster 185
 - used to define a page space 229
 - used to define a path 240
 - nullifying 129
- publications
 - related v
 - required iv
- PURGE parameter
 - in DELETE command 276
 - in EXPORT command 288
 - in IMPORT command 307

Q

- qualified (generic) names 23
 - in ALTER command 72
 - in DEFINE command 46
 - example 216
 - in DELETE command 81
 - in LISTCAT command 75
- QN, print chain option 363

R

- read integrity 30
- read password 27
 - (see also READPW parameter)
- read-only access (see READPW parameter)
- READPW parameter
 - in ALTER command 130
 - in DEFINE command
 - for an alternate index 168
 - for a catalog 260
 - for a cluster 199
 - for a page space 235
 - for a path 245
 - nullifying 130
- record
 - (see also RECORDSIZE parameter)
 - replacement, example of 380
- RECORDS parameter (DEFINE)
 - used to define an alternate index 157
 - used to define a catalog 254
 - used to define a cluster 187
 - used to define a data space 249
 - used to define a page space 229
- RECORDSIZE parameter
 - in ALTER command 131
 - in DEFINE command
 - to define an alternate index 168
 - to define a cluster 199
 - to define a data space 249
- recoverable, making the master catalog 39
- RECOVERABLE parameter (DEFINE) 260

- recovery, catalog 39,93
 - other recovery methods 39
 - the catalog recovery area 93
- recovery area (see catalog recovery area)
- RECOVERY parameter (DEFINE)
 - used to define an alternate index 173
 - used to define a cluster 204
- regaining access to data 40
- regions sharing data 170,201
- RELATE parameter (DEFINE) 61,157
- related publications v
- reloading a catalog 112
- REMOVEVOLUMES parameter (ALTER) 131
 - volume cleanup 38
- removing volumes 131
- removing VSAM ownership of a volume 18,38
- renaming data sets
 - in ALTER command 71,119
 - in IMPORT command 305
- reorganizing data sets 104
- REPLACE parameter (REPRO) 340
- REPLICATE parameter (DEFINE)
 - used to define an alternate index 169
 - used to define a cluster 201
- REPRO command 104,338
 - catalog copying 110
 - catalog unload/reload 112
 - optimizing the performance of 114
 - converting device of user catalog 104
 - example 343
 - format 338
 - parameters 338
 - optional parameters 339
 - required parameters 338
- required publications iv
- requirements, storage
 - VSAM catalog 51
- reserving volumes (see CANDIDATE parameter, VOLUME parameter, and ADDVOLUMES parameter)
- reset condition codes 362
- RESETCAT command 350
 - examples 353
 - format 350
 - function 98
 - JCL Requirements 102
 - multivolume data sets 101
 - parameters 350
 - optional parameter 350
 - required parameter 350
 - reference summary 471
 - requirements 100
 - WORKFILE Space Requirements 101
- respecifying attributes 99,119
- Resource Access Control Facility (RACF) 31
 - when moving entries 82
 - in EXPORTRA 97
 - in ALTER 129
- RESTORE program, IEHDASDR 35
- restoring a cluster's end-of-file values 91
- restoring catalog entries after system failure 93
 - catalog recovery area contents 93
 - listing the 94
 - copying a catalog entry from the catalog recovery area 95
 - examples of 315

- procedure for 93
- resetting catalog entries 98
- restoring the catalog entry that was obtained using the EXPORTRA command 97
- retention date for data set (see also FOR parameter and TO parameter)
- nullifying 130
- overriding (see PURGE parameter)
- RETENTION parameter (ALTER) 130
- reusable data sets 201
- REUSE parameter
 - in DEFINE command
 - used to define an alternate index 170
 - used to define a cluster 201
 - in REPRO command 341
- RN, print chain option 363
- rules of continuation 7

S

- SAM data sets
 - cataloging 223
 - converting 104
 - copying 104
 - deleting 77
 - listing catalog entries for 75
 - printing 114
- SAM to VSAM conversion 104
- SAVRAC parameter
 - in IMPORT command 307
 - in IMPORTRA command 314
- SCRATCH parameter
 - in DEFINE command
 - used to define a generation data group 218
 - in DELETE command 276
 - modified in ALTER command 132
- searching catalogs, order of
 - with the ALTER command 72
 - with the BLDINDEX command 69
 - with the DEFINE command 47
 - with the DELETE command 80
 - with the LISTCAT command 75
- secondary allocation (see CYLINDERS parameter, RECORDS parameter, and TRACKS parameter)
 - amount for the catalog 56
- secondary extents, accessibility to 33
- security
 - authorization routine 31
 - bits in Format-1 DSCB 17
 - passwords 27
- separators 6
- sequence of execution, controlling 5,359
- sequential data set, converting 104
- SEQUENTIALDUMP parameter (LISTCRA) 325
- SET command 362
 - examples 364
 - format 362
- set condition codes 362
- setting up the master catalog 20
- SHAREOPTIONS parameter
 - in ALTER command 132
 - in DEFINE command
 - used to define an alternate

- index 170
 - used to define a cluster 201
- sharing data
 - options modified in ALTER command 132
 - specified in DEFINE command 170,201
- SKIP parameter
 - in PRINT command 330
 - in REPRO command 340
- slashes 5,10
- SN, print chain option 363
- sort parameters (BLDINDEX) 139
- sort workfiles for building an alternate index 67
 - DD statements that describe the 69
 - estimating space requirements for 67
- space allocation for the catalog recovery area (CRA) 57
- space assignment
 - to a catalog 51
 - to a catalog recovery area (CRA) 57
 - to VSAM objects 57
- space estimates
 - for a catalog 51
 - example of 264
 - worksheet for 55
 - for the BLDINDEX sort workfiles 67
 - for the RESETCAT work file 101
- SPACE parameter
 - in DEFINE command 57,249
 - in DELETE command 273
 - in LISTCAT command 319
 - example of output 421
- space, data
 - allocating 57,249
 - defining 57,249
 - deleting 77,273
 - extending 56
 - for a catalog recovery area (CRA) 57
 - space allocation in a catalog's data space 57
 - VSAM objects in a 58
- SPANNED parameter (DEFINE) 203
- spanned records 203
- specifying alternate index information 61
 - descriptive information 62
 - performance options information 62
 - protection and integrity information 62
- specifying cluster information 59
 - descriptive information 59
 - performance options information 60
 - protection and integrity information 60
- SPEED parameter (DEFINE)
 - used to define an alternate index 173
 - used to define a cluster 204
- spreading data sets over volumes (see KEYRANGES parameter)
- STAGE parameter 25
 - in ALTER command 133
 - in DEFINE command
 - used to define an alternate index 173
 - used to define a cluster 204
- staging attributes 25
 - (see also BIND, CYLINDERFAULT, DESTAGEWAIT, NODESTAGEWAIT, and STAGE parameters)
- staging drive 25

- staging from mass storage 25
- starting location
 - in PRINT command 330
 - in REPRO command 340
- STEPCAT catalog 110
 - DD statement for 21,110
 - order of catalog use
 - ALTER 72
 - BLDINDEX 69
 - DEFINE 47
 - DELETE 80
 - LISTCAT 75
- stop testing 363
- stopping location
 - in PRINT command 331
 - in REPRO command 341
- storage requirements
 - VSAM catalog 51
- STRING subparameter (ALTER) 129
- structure of commands 6
- SUBALLOCATION parameter (DEFINE)
 - used to define an alternate index 174
 - used to define a cluster 205
- summary of Access Method Services operations 1
- summary of amendments vii
- SWAP parameter
 - in DEFINE PAGESPACE command 235
 - in DEFINE PAGESPACE example 239
- syntax 5
- system failure, restoring catalog entries after 93
- system's catalogs, example of defining the 366
- SYS1.NUCLEUS, pointer to master catalog in 20

T

- TABLE parameter (PARM) 363
- tape data sets
 - for backup copies 112
 - for transporting data 32
 - open at checkpoint 116
- tape volumes mounted at checkpoint 116
- temporary exportation 83
- TEMPORARY parameter (EXPORT) 289
- terminate processing 360
- terminator, command 8
- TEST parameter (PARM) 363
- THEN clause, in IF-THEN-ELSE command sequence 360
- Time Sharing Option (TSO) 23
 - abbreviated LISTCAT listings 319,436
 - invoking Access Method Services from a terminal 14
- timestamp updating 15
- TN, print chain option 363
- TO parameter
 - in ALTER command 134
 - in DEFINE command
 - used to define an alternate index 174
 - used to define a catalog 261
 - used to define a cluster 205
 - used to define a generation data group 218
 - used to define a nonVSAM data

- set 227
 - used to define a page space 235
 - used to define a path 246
- TOADDRESS parameter
 - in PRINT command 331
 - in REPRO command 341
- TOKEY parameter
 - in PRINT command 331
 - in REPRO command 341
- TONUMBER parameter
 - in PRINT command 331
 - in REPRO command 341
- TRACE parameter (PARM) 363
- tracing outputs 363
- TRACKS parameter (DEFINE)
 - used to define an alternate index 157
 - used to define a catalog 254
 - used to define a cluster 187
 - used to define a data space 249
 - used to define a page space 229
- Translate Table, device type 414
- transporting data sets between systems 32,82
- TSO (Time Sharing Option) 23
 - abbreviated LISTCAT listings 319,436
 - invoking Access Method Services from a terminal 14
- types of commands 4

U

- underlining, in notation conventions vi
- UNINHIBIT parameter (ALTER) 127
- unique data set
 - defining 205
 - space for 45
- UNIQUE parameter (DEFINE)
 - used to define an alternate index 174
 - used to define a cluster 205
- UNIQUEKEY parameter
 - in ALTER command 134
 - in DEFINE command 174
- unload/reload, catalog 112
 - optimizing the performance of 114
- unloading a catalog 112
- UNORDERED parameter (DEFINE)
 - used to define an alternate index 167
 - used to define a cluster 198
- UNORDERED subparameter (IMPORT) 305
- update access 27
- UPDATE parameter
 - in ALTER command 134
 - in DEFINE command 246
- update password 27
 - nonVSAM data sets and the catalog's update password 54
 - (see also UPDATEPW parameter)
- UPDATEPW parameter
 - in ALTER command 135
 - in DEFINE command
 - ALTERNATEINDEX 174
 - catalog 261
 - CLUSTER 206
 - PAGESPACE 236
 - PATH 247
 - USERCATALOG 261

- nullifying 130
- updating a backup copy of a catalog 35
- updating a data set's end-of-file information 91
- updating timestamps 18
- UPGRADE parameter
 - in ALTER command 135
 - in DEFINE command 176
- uppercase, in notation conventions vi
- USAR (see user-security authorization record)
- user catalog
 - (see also master catalog)
 - allocating 51
 - altering 71
 - backing up 35,112
 - calculating space for 51,54
 - cataloging nonVSAM data sets 64
 - connecting to master catalog (IMPORT), 83
 - creating a 253
 - examples 264
 - defined on mass storage volume 51
 - deleting 77
 - disconnecting from master catalog (EXPORT) 83
 - exporting 83
 - importing 82
 - JCL 21,46
 - listing 75
 - order of use
 - ALTER 72
 - BLDINDEX 69
 - DEFINE 47
 - DELETE 80
 - LISTCAT 75
 - preventing duplicate names 46
 - protecting 34
 - relationship to master catalog 20,51
 - reloading 112
 - structure 19
 - transporting 38
 - unloading 112
 - using a backup copy 35
 - volume portability 82
- USERCATALOG parameter
 - in DEFINE command 253
 - in DELETE command 273
 - in LISTCAT command 319
- user's program, invoking Access Method Services from a 13
- user's I/O routines 472
- using generic (qualified) names for cataloged objects 21
- using one object as a model for another object 49
- using passwords to authorize access to data 27
- using the catalog recovery area 93
- user-security-authorization record (USAR)
 - in ALTER command 124
 - in DEFINE command
 - ALTERNATEINDEX 160
 - catalog 254
 - CLUSTER 190
 - PAGESPACE 230
 - PATH 242
 - USERCATALOG 254
- user-security-verification routine (USVR) 31
 - in ALTER command 124

- in DEFINE command
 - ALTERNATEINDEX 160
 - catalog 254
 - CLUSTER 190
 - PAGESPACE 230
 - PATH 242
 - USERCATALOG 254
- USVR (see user-security-verification routine)

V

- VERIFY command 357
 - example 358
 - format 357
 - parameters 357
- volser (see REMOVEVOLUMES, VOLUME, and VOLUMES parameters)
- volume
 - allocating the 58
 - data spaces on a 57
 - dumping a 35
 - identifying the catalog's 54
 - volume cleanup 36,73
 - volume mounting required during DELETE 77
 - volume ownership, VSAM 17
 - removing it 36
- VOLUME parameter
 - in DEFINE command
 - used to define a catalog 254
 - in LISTCAT command 319
 - example of output 419
- volume, mass storage 25
- volume serial number, changing the 37
- VOLUMES parameter (DEFINE)
 - used to define an alternate index 159
 - used to define a catalog 254
 - used to define a cluster 188
 - used to define a data space 251
 - used to define a nonVSAM data set 224
 - used to define a page space 230
- VOLUMES subparameter
 - in IMPORT command 306
 - in IMPORTRA command 312
- volumes, tape, mounted at
 - checkpoint 116
- VSAM catalog (see catalog, master catalog, and user catalog)
- VSAM catalog cleanup 38,275
- VSAM data sets, example of defining 210
- VSAM to SAM conversion 104
- VSAM volume cleanup 36,73
- VSAM volume ownership 17
 - removing it 38
- VTOC (volume table of contents) 17,38

W

- WORKCAT parameter (RESETCAT) 351
- WORKFILE parameter (RESETCAT) 352
- Work file space requirements (RESETCAT) 101
- WORKFILES parameter (BLDINDEX) 139
- write access (see UPDATEPW parameter)

write integrity 31
write operation, verification (see
WRITECHECK parameter)
WRITECHECK parameter
 in ALTER command 135
 in DEFINE command
 ALTERNATEINDEX 176
 CLUSTER 206

USERCATALOG 262

123

3860 Mass Storage System, IBM 25



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM shall have the nonexclusive right, in its discretion, to use and distribute all submitted information, in any form, for any and all purposes, without obligation of any kind to the submitter. Your interest is appreciated.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please 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.*

Please use pressure sensitive or other gummed tape to seal this form.

List TNLs here:

If you have applied any technical newsletters (TNLs) to this book, please list them here:

Last TNL _____

Previous TNL _____

Previous TNL _____

Fold on two lines, tape, and mail. No postage necessary if mailed in the U.S.A. (Elsewhere, any IBM representative will be happy to forward your comments.) Thank you for your cooperation.

Reader's Comment Form

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

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

POSTAGE WILL BE PAID BY ADDRESSEE:

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



Fold and Tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. This form may be used to communicate your views about this publication. They will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

IBM shall have the nonexclusive right, in its discretion, to use and distribute all submitted information, in any form, for any and all purposes, without obligation of any kind to the submitter. Your interest is appreciated.

Note: *Copies of IBM publications are not stocked at the location to which this form is addressed. Please 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.*

Please use pressure sensitive or other gummed tape to seal this form.

List TNLs here:

If you have applied any technical newsletters (TNLs) to this book, please list them here:

Last TNL _____

Previous TNL _____

Previous TNL _____

Fold on two lines, tape, and mail. No postage necessary if mailed in the U.S.A. (Elsewhere, any IBM representative will be happy to forward your comments.) Thank you for your cooperation.

Reader's Comment Form

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

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

POSTAGE WILL BE PAID BY ADDRESSEE:

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



Fold and Tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601